

Exploring the influence of privacy awareness on the Privacy Paradox on smartwatches



Meredydd Williams
Wolfson College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2018

Abstract

While the public claim to be concerned about privacy, they rarely express protective behaviour. This disparity has been labelled the *Privacy Paradox*. The Internet-of-Things (IoT) is ubiquitous and data-collecting. It has therefore been considered a privacy risk. However, the Paradox has not been assessed in this context. Hence, we explored whether the issue could be mitigated in novel environments. Before analysing these devices, we began by confirming the Paradox's prevalence. This was undertaken through street surveys, which suggested that views had little relationship with behaviour. We continued by comparing user perceptions across a range of technologies. Through an online survey [$n = 170$], participants rated 'IoT' devices as most concerning. They were also considered less usable, potentially placing constraints on protection. To dissect the issue, we then conducted detailed interviews. Discussions were undertaken with 40 users, 20 of which had an IoT product. Protection was less common in the IoT, contributing to greater Paradox prevalence. Wearables were particularly prone, and therefore we scoped our focus to smartwatches. Through analysing participant rationale, a lack of awareness appeared the greatest IoT issue. Since educational games had proved effective in the literature, we developed privacy apps. Our online prototype was trialled by 504 smartwatch owners. In our treatment group, we found the Paradox was mitigated in posttest results. Therefore, for our final study, we implemented the first smartwatch privacy game. Through a two-month longitudinal study, we analysed empirical behaviour. Users were given an (Android) Wear OS watch, with half playing the game. Protection increased and persisted in this group, with the Paradox mitigated over an extended period. Our research has novelty, in that this has not been achieved in previous work. Furthermore, by analysing smartwatch behaviour, we provide novel insights into user rationale. As the IoT expands, it is crucial that individuals are informed to make privacy decisions.

Acknowledgements

A DPhil is a long journey comprising of many highs, many lows and many surprises. This thesis would certainly not be possible without the support of several individuals. I am indebted to their kind assistance over the past three years.

Firstly, I would like to thank my two supervisors, Prof Sadie Creese and Dr Jason R C Nurse. Both have offered great feedback and support throughout the DPhil. Sadie has always been helpful, enthusiastic and full of insightful suggestions. Jason provided a valuable red pen and promoted my development as a researcher. On a personal note, they have both been excellent to work with.

Secondly, I would like to express my gratitude to my two doctoral examiners. Prof Fléchais has already provided invaluable advice in the transfer and confirmation stages. I appreciate Prof Briggs travelling down from Northumbria, and look forward to discussing my research.

Thirdly, the moral support has been great within my research team. Jo, Jass and Arnau have all informed my growth as an academic. Their wit and banter has also lightened up the office. My student colleagues have had the greatest influence. It has been a pleasure to study beside Louise, Alastair, Mariam, Mary, Arianna, Marcel, Faisal, AJ, Liz, Rodrigo and Hugo. They appreciate the journey of a doctorate, and I wish them well in their research. Special thanks are also due to Prof Michael Goldsmith, Dr Helena Webb and Lorna Swadling, all of which have supported my work at different stages.

Fourthly, this research would not have been undertaken without the CDT in Cyber Security. I would like to express my gratitude to Prof Andrew Martin, David Hobbs and Maureen York. I am also appreciative of all the support offered by Oxford, Wolfson College and the Department of Computer Science.

Finally, I could not have completed this work without my nearest and dearest. Thank you to my mum for always having confidence in me, even when I don't. And thanks to my partner, Annabel, for tolerating me over the final few months!

Declaration

This thesis is presented in accordance with the regulations for the degree of Doctor of Philosophy. It has been composed by myself and has not been submitted in any previous application for a degree. The research within this thesis was also conducted by myself. Parts of the thesis have been published or submitted as papers, and the list can be found in Section 1.6.

Contents

1	Introduction	1
1.1	Problem Statement	1
1.2	Research Questions	3
1.3	Motivation	4
1.4	Scope	6
1.5	Thesis Outline	8
1.6	Peer-Reviewed Publications	10
1.7	High-Level Progression	13
2	Literature Review	14
2.1	The Concept of Privacy	14
2.2	The Privacy Paradox	16
2.3	The Internet-of-Things (IoT)	26
2.4	Smartwatches	30
2.5	Behaviour Change	34
3	Methodology	46
3.1	Research Philosophy and Approach	46
3.2	Quantitative Techniques	49
3.3	Qualitative Techniques	53
3.4	Threat Models	64
3.5	Consistent Selection Criteria	66
3.6	Ethics	69
4	Can the Privacy Paradox be confirmed in the UK?	71
4.1	Introduction	71
4.2	Methodology	73
4.3	Results and Discussion	82
4.4	Implications	91

5	How do perceptions differ between IoT and other devices?	93
5.1	Introduction	93
5.2	Methodology	95
5.3	Results and Discussion	104
5.4	Implications	112
6	How and why does the Privacy Paradox differ between IoT and other devices?	114
6.1	Introduction	114
6.2	Methodology	115
6.3	Results and Discussion	125
6.4	Implications	141
7	Can we mitigate the short-term prevalence of the Privacy Paradox on smartwatches?	143
7.1	Introduction	143
7.2	Methodology	145
7.3	Gameplay Session Design	160
7.4	Baseline Findings	164
7.5	Gameplay Session Findings	169
7.6	Posttest Findings	171
7.7	Participant Rationale	174
7.8	Implications	178
8	Can we mitigate the medium-term prevalence of the Privacy Paradox on smartwatches?	180
8.1	Introduction	180
8.2	Methodology	182
8.3	Smartwatch Game Design	198
8.4	Baseline Findings	204
8.5	Gameplay Period Findings	213
8.6	Posttest Findings	217
8.7	Participant Rationale	227
8.8	Implications	237

9	Conclusions	238
9.1	Introduction	238
9.2	Summary	238
9.3	Critique	240
9.4	Contributions	247
9.5	Future Work	252
9.6	Concluding Remarks	253
	Appendix A Qualitative Coding Frames	255
A.1	Chapter 5: Product Perception Surveys	256
A.2	Chapter 6: Device Interviews	269
A.3	Chapter 7: Prototype Evaluation	294
A.4	Chapter 8: Longitudinal Study	312
	Appendix B Smartwatch Game Images	348
B.1	Online Prototype Game	348
B.2	Wear OS Game	351
	Appendix C Additional Participant Profiles	356
	Bibliography	357

List of Figures

1.1	High-Level Thesis Progression	13
2.1	Theory of Reasoned Action	35
2.2	Theory of Planned Behaviour	36
2.3	Protection Motivation Theory	37
3.1	Draft Coding Indices	56
3.2	Respondent Validation: Question 8 in Chapter 6	63
4.1	Heatmap: Self-Perceptions vs Online Behaviour	89
4.2	Heatmap: Self-Perceptions vs Empirical Behaviour	89
4.3	Heatmap: Privacy Concern vs Online Behaviour	90
4.4	Heatmap: Privacy Concern vs Empirical Behaviour	91
5.1	Device Category Selection	102
5.2	Device Mean Factor Ratings	108
6.1	Qualitative Concern Scaling	130
6.2	Heatmap: Concern Scores vs Behaviour Scores	139
7.1	Experimental Structure	151
7.2	Introductory Video	160
7.3	Prototype Educational Game	161
7.4	Heatmap: Concern Scores vs Behaviour Scores	168
7.5	Comparison of Privacy Paradox Prevalence	174
8.1	Experimental Structure	187
8.2	Smartwatch Privacy Game	198
8.3	Pretest Smartwatch Privacy Concerns	209
8.4	Privacy Game Challenge Completion Rates	215
8.5	Privacy Paradox Changes: Treatment Group	228
8.6	Privacy Paradox Changes: Control Group	228

List of Tables

3.1	Example Coding Frame: Question 12 in Chapter 8	57
3.2	Example Framework Matrix: Chapter 8 Gameplay Responses	60
4.1	Street Survey Questionnaire	77
5.1	Product Perception Questionnaire	98
6.1	Semi-Structured Interview Questions	120
6.2	Example Coding Frame: Data Deletion Incident	129
6.3	Example Coding Frame: Password Usage	130
7.1	Prototype Study: Concern Questionnaire	154
7.2	Pretest-Posttest Scores	171
8.1	Longitudinal Study: Concern Questionnaire	193
8.2	Game Evaluation Questionnaire	195
8.3	Posttest Interview Questions	196
8.4	Privacy Game Challenges	201
8.5	Participant Changes in Protective Behaviour	222
8.6	Longitudinal Privacy-Protective Behaviour	225

Chapter 1

Introduction

We will introduce the topic of the thesis in this initial chapter. Previous research has demonstrated that while users claim to care about privacy, they often do little to protect themselves [47, 153, 214]. This disparity¹ between concern and behaviour has been labelled the *Privacy Paradox* [279]. Although the topic has been studied extensively in both online and offline environments [279, 290], the Internet-of-Things (IoT) is underexplored. Since smart devices can pose privacy risks [126, 181, 300], it is important that users are supported.

Overview. In this chapter, we will begin by outlining the problem we seek to address. Then we will discuss our central research question, which builds on this issue. Through the subquestions, we deconstruct how the topic is approached. This will be followed by motivation, supporting both our research theme and the studies conducted. We will then scope the focus of our thesis, highlighting the fields we seek to inform. We continue by describing the general structure of the document. Penultimately, we will highlight the works published during the doctorate. Finally, we illustrate how our chapters synthesise to address our central research question.

1.1 Problem Statement

The Privacy Paradox. The public frequently claim concern for their privacy [167, 304]. However, individuals rarely take action to protect personal data. Users often fail to read policies [155], check permissions [128] or safeguard their information [270]. They also use services that can endanger their privacy, such as social networking sites [57]. Even when an intent is expressed to act privately, it is frequently followed by lax actions [279]. We define the Paradox as the “*discrepancy between the expressed*

¹‘Disparity’ is used in this thesis as a shorthand for the disparity between concerns and behaviour.

concern and the actual behavior of users". This originates from Barth and de Jong's systematic literature review of the topic [47]. Despite extensive research in this area, the Paradox has never been mitigated² for an extended period [194].

It is unclear whether the issue is truly 'paradoxical' or whether the term was chosen for alliterative rhetoric. Some claim it could be likened to hypocrisy [238], while others argue that decisions are rational [365]. We outline and justify our own interpretation within the Literature Review (Section 2.2). Regardless, behaviour appears inconsistent with privacy concerns. These actions can place users at risk, particularly since consequences are often not understood [194]. Furthermore, when individuals reflect on their decisions, they frequently express regret [385]. We wish to support users in making informed privacy decisions.

Internet-of-Things. The IoT describes a nebulous and varied collection of 'smart devices' [377]. These products can range from wearables to appliances, drones to home automation. Such technologies offer several benefits to society, including industrial productivity [231] and patient-led healthcare [283]. However, research has highlighted the potential risks to privacy [181]. Products, and their manufacturers, often rely on data collection for revenue [327]. With these devices pervading our homes, this can be done surreptitiously. The IoT is nebulous and heterogeneous [345], with a variety of products from a variety of vendors. This magnifies the unfamiliarity of the novel environment [42]. Their usability has also been criticised in several studies [90, 143, 183, 380]. While these issues might constrain protective behaviour [158, 230, 247, 317], the Paradox has never been analysed in this environment.

Smartwatches. As our research progressed, we scoped to the products which most-exhibited the Paradox. Through analyses of concern and behaviour, the issue appeared common to wearable devices. These gadgets are often small and lacking usability [52]. They can also house a variety of sensitive data, ranging from text messages to phone contacts [118]. When questioned on wearable risks, individuals express concern [232]. However, they appear to overlook their protective settings [371]. As smartwatches grow in popularity, personal data might be placed at risk. To support user privacy, we wish to mitigate the Paradox's prevalence³ in this context.

²The term 'mitigate' is frequently used. We do not claim to eliminate the Paradox, but make it "*less severe, serious, or painful*". www.oxforddictionaries.com/definition/mitigate

³'Prevalence' is also commonly used. It is used in the sense of "*the degree to which something is prevalent*", as defined in <https://www.merriam-webster.com/dictionary/prevalence>. Indeed, we do not claim that the Paradox is widespread or 'prevalent' in all cases.

1.2 Research Questions

Before seeking to mitigate the issue, we must understand behavioural rationale and mitigative approaches. Therefore, we conduct a range of analyses.

Our central research question can be defined as: *Can the Privacy Paradox be mitigated in the context of smartwatches?* This query was selected for several reasons. We believe that the IoT could limit privacy protection. This should be particularly common in smartwatches, where data is collected [118] but interfaces are constrained [183]. Since concerns are often based on principle [278], we deem them less likely to change. Therefore, without mitigative approaches, the Paradox might increase in prevalence. To empower device users, we aim to provide education and practice. Through supporting informed decisions, behaviour may realign with concern. This should reduce inadvertent risks [194] and regrettable actions [385].

The central research question is broken down into several subquestions. These are each addressed in our subsequent chapters. They are justified below:

SQ1. Can the Privacy Paradox be confirmed in the UK?

This is addressed in the study documented in Chapter 4. Before exploring the IoT and smartwatches, it was important to confirm the presence of the Paradox. Although it has been extensively analysed in the US [4, 122, 148, 163], research is sparse in the UK [68, 397]. With privacy perceptions varying by culture [18], it was prudent to start in this manner.

SQ2. How do perceptions differ between IoT and other devices?

This is addressed in the study documented in Chapter 5. Since IoT risks are often highlighted in the media [131, 154, 273], we believe consumers might be concerned. If these products also lack usability and familiarity, they may constrain protection [158, 230, 317]. Such a result could contribute to a prevalent Paradox. To explore perceptions of several devices, we conduct an online survey.

SQ3. How does the prevalence of the Privacy Paradox differ between IoT and other devices?

This is addressed in the study documented in Chapter 6. As expressed earlier, we believe the IoT could increase the Paradox. To assess the prevalence of the issue, concerns and behaviour must be explored across different environments. By using older technologies as a baseline, we can analyse how the IoT differs. This is achieved by studying the matter through semi-structured interviews.

SQ4. Which factors contribute to the Privacy Paradox in the IoT?

This is also addressed in the study in Chapter 6. To support protection, we do not seek to merely confirm the Paradox. We must also understand the factors that influence the issue. Once these components are known, we can design mitigative approaches. Therefore, it was essential to explore the rationale of IoT users. This was undertaken through the semi-structured interviews. To refine our approach, factors were further explored in Chapters 7 and 8.

SQ5. Can we mitigate the short-term prevalence of the Privacy Paradox on smartwatches?

This is addressed in the study documented in Chapter 7. In Chapter 6, we found that wearables are particularly prone to the Paradox. Since we wish to support protection, we therefore scoped our focus to smartwatches. Based on the influential factors, we could develop approaches to mitigate the Paradox. Through evaluating our success over the short-term, we refined our design. At this stage, we implemented a prototype ‘educational game’ [51].

SQ6. Can we mitigate the medium-term prevalence of the Privacy Paradox on smartwatches?

This is addressed in the study documented in Chapter 8. To answer our central research question, we sought to mitigate the prevalence of the Paradox. Although this had been achieved over a brief period (Chapter 7), behaviour might relapse when privacy loses salience [189]. Therefore, we studied the matter over the medium-term⁴. Techniques were refined based on results from our prototype. Through a longitudinal study⁵ with the first smartwatch privacy game, we addressed our central research question.

1.3 Motivation

We now move forward to discuss our motivation in pursuing the research questions.

⁴The ‘medium-term’ is “*the period of time which lasts a few months or years beyond the present time*”. In absence of a preferable description, this reflects the length of our two-month study. <https://www.collinsdictionary.com/dictionary/english/medium-term>

⁵Since this work employs “*repeated measures to follow particular individuals over prolonged periods of time*”, it is regarded as a longitudinal study [80]. It has similarity to other privacy research which defines itself in this manner [384]. It could be alternatively termed a ‘field study’.

Empowerment. The Paradox has been confirmed in a variety of environments [279, 290]. Whereas individuals might value the principle of privacy [278], they often neglect to use protection [4]. This can be a conscious decision, with individuals balancing privacy against functionality [117]. However, research suggests that a lack of awareness is often influential [114]. In these cases, the benefits of disclosure tend to be overestimated [157]. Furthermore, users might assume they are safe and reveal information unwittingly [194]. When reflecting on their actions, they then often express regret [385]. We wish to empower individuals to protect their privacy. Through implementing educational games, we seek to highlight the potential consequences. Even if a person then avoids protection, at least they make an informed choice.

Market forces. Privacy is commonly regarded as a secondary goal [188]. When considering modern technology, other forces appear more influential. Companies seek to release their products quickly in order to corner the market [334]. Functionality is often prized above protection, since it is more immediately salient [334]. Even when privacy features are included, vendors have little incentive to make them usable [59]. Many products are funded by data collection [334], and it is known that defaults are seldom changed [244]. In such environments, protection is unlikely to be used. For consumers to have a chance against vendors, their privacy must be supported.

The Internet-of-Things. The IoT is expanding rapidly, with the market predicted to reach \$267b by 2020 [98]. While these products offer great convenience, their privacy threats are well-documented [181]. They can collect large quantities of data through completing their tasks [300]. Smart TVs have been found capturing audio, before transmitting the content without encryption [154]. Home automation systems can also monitor their environment, enabling the inference of personal details [376]. IoT products are unfamiliar [42] and often lacking in usability [90, 143, 380]. With consumers constrained in this manner, smart devices might pose a threat to privacy [158, 230, 317]. Unless users are supported, their data might be placed at risk.

Smartwatches. Smartwatches are valued for their wearable functionality and quick convenience. Due to their appeal, user studies have begun to explore these devices [196, 306]. However, privacy/security behaviour has not been empirically analysed. We find this surprising, especially since personal data can be stored [118]. Due to this research lacuna, we were keen to discover how settings are used. Smartwatches have familiar features, including screen locks and app permissions. But since settings are neglected on phones [128], protection might be rare. Through studying empirical behaviour over several months (Chapter 8), we could assess how users behave.

Understanding. The Paradox has been analysed through many prior studies. However, to address the issue in a novel environment, it first must be understood. For this reason, we were motivated to collect rich qualitative data. This was undertaken through in-depth interviews in Chapters 6 and 8. We also surveyed the views of 504 smartwatch owners in Chapter 7. This granted us an opportunity to understand why privacy might be neglected. Based on participant rationale, we can support the development of improved protection. In this manner, our research is informed through interviews and qualitative data.

1.4 Scope

Since research concerns a defined topic, scoping is essential. Such an approach highlights which areas are under consideration and which remain for future exploration.

Information privacy. Privacy is a nebulous principle [1], encompassing confidentiality, solitude, and other facets [352]. An individual might want their data to be private, but be happy for their identity to be known. Privacy is also inherently subjective [289], cultural [18] and contextual [278]. For these reasons, it was crucial we scoped our research area. Clarke [92] divided the topic into four elements: information privacy, media privacy, interception privacy and bodily privacy. Since we wish to explore technological interactions, we constrain our focus to information privacy. He defined this as “*the interest an individual has in controlling, or at least significantly influencing, the handling of data about themselves*” [92]. Hence, we compare concerns of information privacy against information privacy behaviour.

Consumer devices. As previously mentioned, the IoT is nebulous and heterogeneous [345]. The environment is defined as a “*global network interconnecting smart objects*” [268], which itself encompasses a range of technology. However, since we wish to analyse consumers, our selection requires constraint. Most individuals are unlikely to interact with industrial devices, such as factory control systems. Furthermore, these technologies are expensive, rare and not amenable to user studies. Therefore, we scoped our research to consumer devices. These products are more likely to be owned by our participants, and hence we can solicit informed opinions.

Smartwatches. As our research progressed, we evaluated a narrower set of products. This was important, as it would be infeasible to study all consumer devices. To explore whether the prevalence of the Paradox could be mitigated, we selected the most-prone environment. Through Chapter 5 and 6 analyses, wearables appeared to pose the greatest issues. Of these devices, smartwatches were deemed appropriate for

study. They are functional, popular and possess a range of privacy settings. Through studying their usage, we could assess the frequency of protective behaviour. Furthermore, smartwatches are amenable to modification and analysis. This provided an opportunity to implement our educational games. By exploring user actions, we explored whether the Paradox could be mitigated.

UK-centric. While the Paradox has been explored extensively in the US [4, 122, 148, 163], research is sparse in the UK [68, 397]. Brown [68] undertook early work, but his 2001 study sampled only 12 participants. Although Zafeiropoulou et al. [397] performed a larger analysis, over half their respondents were non-UK. Since the topic is cultural [18], perceptions might differ across different nations. For example, privacy protection varies in US and EU law [336]. Although the constitution provides safeguards in the US, the same does not apply in Britain [132]. This might lead to different privacy expectations, potentially impacting concern and behaviour.

Our street survey, documented in Chapter 4, is conducted across several British locations. The product perceptions survey (Chapter 5), was advertised on local and national messaging boards. Since these boards were UK-specific, it is unlikely that many participants were outside the country. For our semi-structured interviews (Chapter 6), respondents were recruited from the Oxford public. We then trialled our prototype game through an online study (Chapter 7). To receive participants, we made use of a UK crowdsourcing platform (Prolific). The user pool is predominantly British⁶, and Amazon Mechanical Turk is more popular in other states⁷. Finally, our longitudinal study participants (Chapter 8) originated from several nations. However, they were resident to this country and the procedure took place locally. Therefore, we analyse behaviour within a UK environment. Our research makes no claim to be representative of the entire nation. However, the findings may apply to a subset of society. If the Paradox is consistently confirmed, the issue might be prevalent.

Awareness. As will be outlined, we seek to mitigate the Paradox by providing awareness⁸ and education. This was informed by both the literature [114, 309] and rationale from our Chapter 6 interviews. Since the prevalence of the issue appeared to decrease, we are pleased with our approach. However, we accept that many factors influence the Paradox. A person might be aware of risks, but care more about functionality [103]. Alternatively, their decisions might be constrained by cognitive

⁶<https://www.prolific.ac/demographics/>

⁷<http://demographics.mturk-tracker.com/#/countries/all>

⁸We use ‘lack of awareness’ as a convenient shorthand to denote a low amount of privacy recognition. We do not wish to imply that any individual has zero knowledge of a matter.

biases [6]. Although we targeted awareness, the Paradox might be addressed through other means. We encourage future work to investigate these avenues.

Behaviour. To mitigate the issue, we sought to encourage protective behaviour. With the Paradox being a disparity between actions and concerns, we could have targeted the latter point. We decided against this for several reasons. Firstly, since opinions are often principled [278], they might be challenging to adjust. Secondly, it would be improper to convince individuals that they face little risk. Finally, we wished to equip users with useful skills. Since our approach appeared successful, we are confident in our choice.

1.5 Thesis Outline

To offer an overview of the thesis, we will now outline the contents of each chapter. This summary will highlight the research undertaken and our coherent narrative.

2. Literature Review

Chapter 2 contextualises our research in the existing literature. This serves four main purposes: to introduce the reader to the general theme; to demonstrate a sufficient knowledge of the area; to present which topics were previously unexplored; and to highlight how our research fills this lacuna. We begin by discussing both the principle of privacy and the Privacy Paradox. The IoT is considered next, alongside the risks that it might pose. Since smartwatches are studied in later chapters, these products are then outlined. Once previous studies have been highlighted, we describe behaviour change techniques. Finally, we present an overview of educational games.

3. Methodology

Our five studies sought to address our central research question. To achieve this, the analyses required alignment in a consistent direction. In Chapter 3, we highlight how this is undertaken. For brevity, we also describe techniques which are used frequently. We begin by discussing high-level details, such as research philosophy and methodological approach. We then justify our selection of quantitative techniques. Our qualitative analyses are also outlined in detail, including data validation. Since we wish to gauge concern in a grounded manner, we required feasible threat models. These are discussed, alongside criteria for selecting concern contextualised questions and protective tools. Finally, as we sought to act responsibly, we explain our ethical practices.

4. Can the Privacy Paradox be confirmed in the UK?

Chapter 4 introduces our first study. Through this work, we sought to confirm the Paradox within the UK. We conducted street surveys at four sites across the nation. Through our 112 responses, we found claimed concern was high. However, protection was inconsistent and disclosure was common. Since the Paradox appeared prevalent, we could continue by exploring the IoT. To assess the influence of privacy awareness, we also recruited security researchers. They took greater action to protect their data. This provided an early indication that education might be persuasive.

5. How do perceptions differ between IoT and other devices?

In Chapter 5, we discuss our comparative study. This assessed perceptions across a range of consumer technologies. We recruited 170 participants through an online survey. Individuals evaluated six devices: three IoT and three deemed to be less-novel. In addition to privacy, ratings were given on usability, familiarity and utility. When comparing products, the IoT provoked the greatest concern. It was also considered less usable and familiar, potentially constraining protection [158, 230, 317]. Since wearables were rated particularly poorly, this supported their later selection.

6. How and why does the Privacy Paradox differ between IoT and other devices?

We explored concerns and behaviour in detail in Chapter 6. Through comparing a range of devices, the Paradox was examined in the IoT. To achieve this, we conducted semi-structured interviews with 40 device owners. Concerns appeared strong regardless of the product. However, protection was rarely used on IoT technologies. As asserted, the Paradox was more prevalent in this environment, particularly wearables. Our interviews also enabled an exploration of behavioural rationale. Since a lack of awareness appeared the largest factor, we chose to design educational approaches.

7. Can we mitigate the short-term prevalence of the Privacy Paradox on smartwatches?

In Chapter 7, we describe the development and evaluation of our prototype game. Wearables appeared prone to the Paradox, and hence we scoped our focus to smartwatches. Since serious games are effective in providing education [394], we designed a prototype application. This online game was trialled and refined by 504 smartwatch owners. As before, concerns were strong in both our treatment and control groups.

However, only the former increased their protection. This appeared to reduce the Paradox. With the issue mitigated, we were prepared for a longer-term analysis.

8. Can we mitigate the medium-term prevalence of the Privacy Paradox on smartwatches?

Chapter 8 outlines our longitudinal study. Based on feedback from our prototype, we refined an educational game. It was extended and then implemented as an (Android) Wear OS app. 10 participants were given smartwatches for a two-month period. Concerns were surveyed in pretest and posttest through brief questionnaires. After 18 days, the treatment group received the privacy game. The control group were given a similar non-privacy app, reducing confounding variables [344]. While former individuals began to use protection, latter participants took little action. Concerns became slightly moderated in the treatment group, further mitigating the Paradox. Through this, we addressed our central research question. The study was concluded by semi-structured interviews, enabling the extraction of rationale. Users were supported in making informed decisions, even though some disclosed data for functionality.

9. Conclusions

Chapter 9 ties together the narrative from the preceding studies. After reviewing the content, we demonstrate the answering of our research questions. We then include a critique of the conducted work. The chapter continues by outlining our novel and significant contributions. Penultimately, we discuss future work based on these findings. We conclude with a high-level overview.

1.6 Peer-Reviewed Publications

Through our research, we have published several peer-reviewed papers. These publications both testify to the quality of the content and the development of the researcher. In all but Paper 3, the doctoral student conducted the research. For Paper 3, the student collaborated with a colleague, with the former drafting the proposed ‘future scenarios’. Since that work does not form part of this document, all thesis research was undertaken by the doctoral student.

1. **M. Williams**, A study of society’s perception of online privacy. *Proceedings of the 1st Interdisciplinary Cyber Research Workshop* pp. 28-29, 2015.

In this initial abstract, we proposed the street survey contained in Chapter 4. We also designed the methodology and theorised on expected results. The paper hypothesised that there would be a contrast between concern and behaviour. This was confirmed in the survey, suggesting the existence of the Paradox.

2. **M. Williams** and J. R. C. Nurse. Optional data disclosure and the online privacy paradox: A UK perspective. *Human Aspects of Information Security, Privacy and Trust in Lecture Notes in Computer Science* 9750 p. 186-197, 2016.

This paper describes the methodology, results and conclusion from our privacy street survey. Through this, it highlights the existence of the Paradox. This work both contributes to Chapter 4 and motivates the subsequent research.

3. **M. Williams**, L. Axon, J. R. C. Nurse and S. Creese. Future scenarios and challenges for security and privacy. *Proceedings of the 2nd International Forum on Research and Technologies for Society and Industry*, 2016.

While this paper did not form a chapter, it considered the influence of the IoT on privacy. We conducted a literature review of ‘future challenge’ documents produced by governments and academia. We then outlined 10 potential scenarios, including the growth of the IoT and the reinterpretation of privacy. These findings informed our motivation and general rationale.

4. **M. Williams**, J. R. C. Nurse and S. Creese. The perfect storm: The privacy paradox and the Internet-of-Things. *Proceedings of the 11th International Conference on Availability, Reliability and Security (ARES)* pp. 644-652, 2016.

This position paper detailed our general premise. It introduced the Paradox and outlined the factors considered to be contributory. The work went on to describe the IoT and the idiosyncrasies of this environment. We then synthesised the components, positing that the Paradox would be exacerbated. This work contributed to Chapter 2, and the motivation behind Chapters 5 and 6.

5. **M. Williams** and J. R. C. Nurse. Perspectives on privacy in the use of online systems. *Proceedings of the 30th British HCI Conference*, 2016.

To assess behaviour fairly, protective approaches should be feasible. Therefore, we explored which tools were known to non-expert users. We first asked 35 students what they defined to represent protective behaviour. This was compared against the techniques mentioned in 35 well-cited privacy studies. As expected,

complex approaches were often omitted by our sample. When selecting tools in Chapters 6, 7 and 8, these approaches were not included.

6. **M. Williams**, J. R. C. Nurse and S. Creese. Privacy salience: Taxonomies and research opportunities. *IFIP Advances in Information and Communication Technology* 498 pp. 263-278, 2017.

Since a lack of awareness was common, we decided it was wise to survey privacy salience [189]. We conducted a systematic literature review, identifying possible research lacunas. IoT lab/field studies appeared underexplored, and this was addressed in Chapters 7 and 8. There was also an opportunity for ‘Salience-Enhancing Technologies’, with this encouraging our educational games. As outlined, this publication informed the progression of our research.

7. **M. Williams**, J. R. C. Nurse and S. Creese. “Privacy is the boring bit”: User perceptions and behaviour in the Internet-of-Things. *Proceedings of the 15th International Conference on Privacy, Security and Trust (PST)*, 2017.

This paper describes the research conducted in Chapters 5 and 6. It first outlines the methodology and results of our comparative survey. After discussing the findings, it introduces the semi-structured interviews. Through comparing a range of devices, it demonstrates the prevalence of the Paradox. This work was initially displayed as a poster at the 13th Symposium on Usable Privacy and Security (SOUPS2017). It was then published as a full paper at PST2017.

8. **M. Williams**, J. R. C. Nurse and S. Creese. (Smart)watch out! Encouraging privacy-protective behavior through interactive games. *International Journal of Human-Computer Studies*, 2018 (under review).

This submission describes our Chapter 7 prototype study. After outlining game design, it highlights how the Paradox appears to be prevalent. The paper continues by detailing posttest results, presenting that the issue is mitigated. It concludes by considering user rationale, informing the Chapter 8 refinements.

9. **M. Williams**, J. R. C. Nurse and S. Creese. Smartwatch games: Encouraging privacy-protective behaviour in a longitudinal study. *Computers in Human Behaviour*, 2018 (under review).

In our final submission, we outlined the longitudinal study. The paper first justifies the methodology and describes the educational techniques. We then enumerated pretest, gameplay and posttest results. As previously mentioned, it

suggested that the Paradox can be mitigated over an extended period. As this addressed our central research question, it provided the content for Chapter 8.

1.7 High-Level Progression

Our research follows a coherent narrative, with each study informing subsequent analyses. We believe this is highlighted in the Introduction. To illustrate the relationship, a graphic is provided in Figure 1.1. Research subquestions are in red, findings are in green, and literature-inspired considerations are in yellow. As shown, each chapter builds on the findings of its predecessors. In synthesis, these works addressed the central question: *Can the Privacy Paradox be mitigated in the context of smartwatches?*

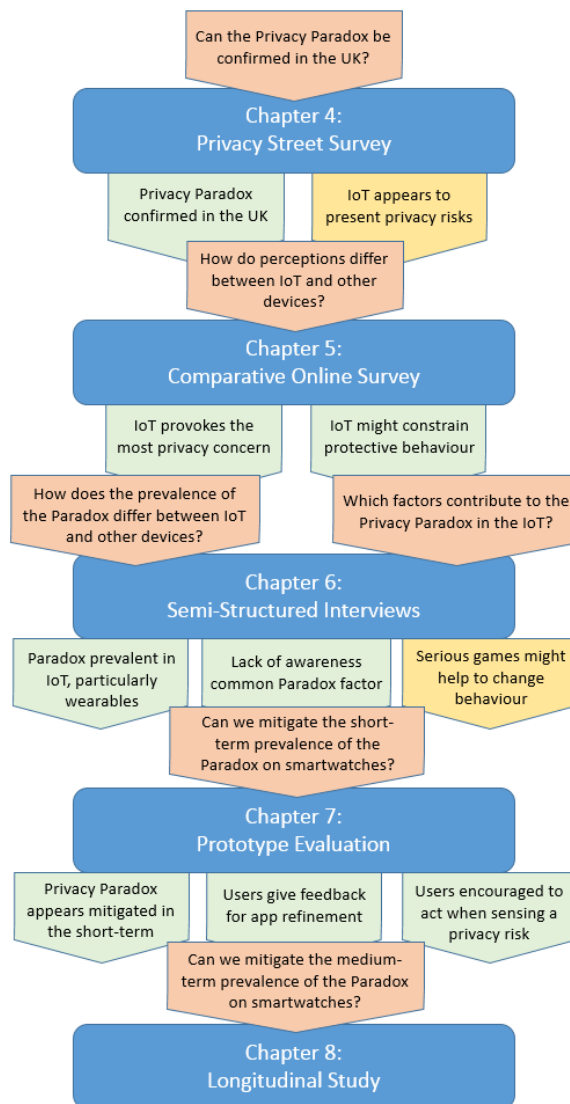


Figure 1.1: High-Level Thesis Progression

Chapter 2

Literature Review

In this chapter, we outline the literature relevant to our research. We also highlight underexplored areas and provide context for our studies. The concept of privacy is discussed first, before we introduce the Privacy Paradox. After discussing the range of interpretations, we turn our attention to the Internet-of-Things. Since we later scope to smartwatches, this environment is also outlined. Finally, we discuss the process of behaviour change and the approaches that we considered.

2.1 The Concept of Privacy

General concept

Historical. Privacy has been an important concept throughout human history, with great civilisations and philosophers considering the topic. The Code of Hammurabi, enshrining Ancient Babylonian law, protected the home against intrusion [109]. Article 21 states that if “*a man makes a breach into a house, one shall kill him in front of the breach and bury him in it*” [228], reflecting that solitude was valued in the Bronze Age. Since the fourteenth century, a particular ‘right to privacy’ was enforced in England. This resulted in individuals being sent to court for eavesdropping [182]. The privacy abuses of King George III were among the reasons the US seceded from Britain [95]. The resultant Bill of Rights enshrined privacy in several amendments, including the first (freedom of speech), the fourth (freedom from unreasonable search and seizure) and the fifth (freedom from self-incrimination) [352].

Modern. Warren and Brandeis [386] placed privacy in the modern democratic consciousness through their 1890 work ‘The right to privacy’. The pair defined a ‘right to be let alone’, claiming individuals possess an “*inviolable personality*” in the face of media surveillance. Privacy is now enshrined as both a legal and human

right in many nations across the world. Article 17 of the International Covenant on Civil and Political Rights states that “*no one should be subjected to arbitrary or unlawful interference with his privacy*” [326]. European privacy has been recently supported through the General Data Protection Regulation (GDPR). This legislation gave greater rights to citizens, while placing additional responsibilities on companies. Following the recent Facebook controversy [156], it appears that privacy is in vogue.

Facets of privacy

Context. Privacy is complex and nebulous topic, incorporating a wide variety of different facets. It can relate to several concepts, including solitude, autonomy and confidentiality [352]. It is also contextual in nature [278], reflecting that perceptions can vary across different situations. For example, a person’s expectation of privacy might differ in a police station to a public street. Therefore, opinions are likely to be inconsistent across a range of digital environments. For this reason, it is crucial that considerations are scoped to particular technologies. In our work, concerns and behaviour are contextualised around specific products.

Culture. Privacy is also subjective [289], with individuals evaluating their data in different respects. While one person might disclose their information readily, another may oppose online access. Furthermore, perceptions can differ greatly based on culture and region [18]. There is a strong social norm for privacy within the western world [352]. However, viewpoints are even found to vary between the US and European states [387]. Privacy expectations are often collectivist in Asia, though this also differs by country [186]. We did not wish to influence our analyses through cultural comparisons. Therefore, our studies were conducted in the UK with UK residents.

Scope. Due to its complexity, privacy is challenging to define [74]. To aid analysis, it has been deconstructed into several concepts. Burgoon [74] found distinctions between four elements: informational privacy, social privacy, psychological privacy and physical privacy. While a person might value their personal space, they might not object to data disclosure. Clarke [92] undertook a similar task, enumerating information privacy, media privacy, interception privacy, and bodily privacy. Since we wished to explore technological interactions, many concepts are out of scope. We studied ‘information privacy’, defined as “*the interest an individual has in controlling, or at least significantly influencing, the handling of data about themselves*” [92].

2.2 The Privacy Paradox

Concept. We now introduce the central topic, the Privacy Paradox [279]. Individuals frequently claim concern for their privacy. A 2018 poll [167] found that 78% believe it extremely important that their data is private. Since 75% would not buy from an invasive company, individuals appear concerned. However, we rarely take action to protect our data. We often fail to read policies [155], check permissions [128] or safeguard our information [270]. This suggests that a disparity is present between claimed concerns and behaviour. Although many have explored this matter (as highlighted below), the Paradox has never been mitigated over a period of time. Our research seeks to fill this gap, since users might endanger their data [194] or regret their actions [385].

Studies. The Privacy Paradox was first highlighted by Brown [68] through his qualitative interviews on Internet use. He discussed interactions with 12 individuals from a variety of backgrounds. A disparity was initially found when considering the role of loyalty cards. Even participants with strong concerns would reveal information through these schemes. Users also disclosed their details to online portals, despite claiming to distrust the degree of data collection. This contrast between concern and behaviour was deemed to be “*something of a privacy paradox*” [68].

Spiekermann et al. [356] also conducted early work on the topic. They studied 171 users in the context of online shopping. Participants were first given a questionnaire to evaluate their opinions. Most claimed to value their information privacy. Users then undertook a simulated shopping task. Despite their concerns, sensitive details were frequently disclosed. This suggested a disconnect between claims and actions.

Social networking sites provided another environment to study the issue. Acquisti and Gross [4] explored concerns and interactions on a nascent Facebook. This was undertaken through an online survey with 294 students. Even when participants were worried, they were found to join the network. They also shared large quantities of personal information. Since understanding of the platform was low, a lack of awareness was likely influential.

At this point, the phenomenon had gained some credibility. However, it was popularised by the empirical work of Norberg et al. [279]. They conducted two studies, with the first comparing claimed preferences with disclosure behaviour. Participants were initially questioned on their willingness to reveal data. 12 weeks later, the same information was requested from the same individuals. Regardless of the sensitivity, disclosure was far greater than what had been claimed. An adapted approach was

used in the second study, with similar results. These findings confirmed the existence of a Privacy Paradox. However, other work sought to understand the disparity.

Privacy calculus

Concept. The Paradox has been justified through a range of constructs. The most intuitive is that of ‘privacy calculus’, which weighs the risks and rewards of technology [107]. These calculations assume the presence of ‘homo oeconomicus’, the economic individual [211]. This is defined as a consumer who makes decisions to maximise their benefit. A person may perceive risks, but deem functionality to be more important. While they have concerns, they place greater weight on other factors.

The ‘calculus of behaviour’ was originally introduced by Laufer and Wolfe in their 1977 work [229]. However, at this point it was not applied to technological interactions. The concept was revived by Culnan and Armstrong [107], who considered the matter in the modern world. They explained that a privacy calculus is undertaken before consumers purchase goods. If the user perceives the vendor to be fair, they are more likely to disclose their information. The model was extended for online interactions by Dinev and Hart [117]. They validated the construct through a survey with 369 participants. Based on Structured Equation Modelling (SEM), it appeared that concerns do inhibit online interactions. However, the advantages of the Internet often outweighed this trepidation.

Limitations. When adopting privacy calculus, individuals are assumed to act in a rational manner. Whereas traditional economics assumes this, it is challenged by a growing body of research. Since the 1970s, the role of heuristics and biases have been explored. Behavioural economics demonstrates that decision-making is affected by many factors [76]. When individuals have less knowledge, they have been found to overestimate the benefits [157]. Furthermore, without a degree of understanding, risks cannot be accurately perceived [350]. Since calculus comparisons are constrained, we explore the influence of other factors.

Biases, heuristics and bounded rationality

Bounded rationality. Individuals are not deemed to act ‘irrationally’, but to have bounds on their decision-making [347]. In a perfect world, we might have infinite information, infinite cognition and no predilections. Unfortunately, these assumptions do not hold for the human mind. This can have significant implications for privacy

judgements. For example, a person might be concerned about online shopping. However, if their purchase is time-pressured, they might overlook the threats.

Acquisti and Grossklags [6] explored this issue in their pioneering work. They described how privacy decisions are unlikely to be perfectly rational. Firstly, users will probably be reliant on incomplete information. Therefore, they lack awareness of all the risks and rewards. Secondly, even if information is complete, it is infeasible to process all these details. Such an act would be both mentally-challenging and prohibitively time-consuming. Finally, even if the above constraints were removed, users would be prone to cognitive biases. For example, they might seek the immediate gratification of using a service [17]. Veltri and Ivchenko [378] explored the influence of cognitive scarcity. When their participants became fatigued, they were found to disclose more information. This provides another partial explanation of the Paradox. People can express concerns easily, but then be impeded when undertaking actions.

Availability heuristic. A heuristic is a quick but suboptimal approach to solving a problem. Through availability bias, salient events are considered more likely [370]. For example, if an aeroplane crash is recent, people might overestimate the risk. Similarly, incidents might be judged improbable if they have not been experienced. Hallam and Zanella [163] studied the matter in terms of privacy incidents. They examined the Paradox through Construal Level Theory, which studies whether concepts are considered abstract or concrete. It was found that concerns had no direct influence on behaviour. Since privacy risks were not salient, they were perceived as being abstract and less likely.

Hyperbolic discounting. This concept states that individuals would prefer an immediate small reward to a delayed larger reward [225]. If an individual discloses online, they might face a long-term risk. However, since the benefit is instant, they disregard their concerns. The concept was introduced to privacy by Acquisti and Grossklags [5], but an experiment was not conducted. Alashoor and Baskerville explored a related topic; that of cognitive absorption [17]. This concerns situations where individuals are engrossed in software. The researchers presented three privacy vignettes, with each highlighting the risks of gratification. They concluded that the pleasure of interaction can cause threats to be underestimated.

Optimistic bias. Even when users express concerns, they might consider themselves free from risk. In these cases, they can lack motivation to protect their data. Optimistic bias describes how people judge themselves to be less susceptible than others [170]. Cho et al. [87] analysed perceived vulnerability through a 910-person survey. Participants were found to distinguish between two levels of privacy risk: the

personal level and the societal level. Whereas they believed others were prone, they were optimistic about their own chances. This might justify why individuals avoid protection, even when they possess concerns.

Positive affect. In privacy calculus, the rewards and risks should be compared rationally. However, emotion and affect have been shown to be influential. If a person has a positive mood, they might be more likely to share data [202]. The matter was explored by Kehr et al. [202], who recruited 480 participants. Individuals took part in an online experiment, interacting with one of two screens. One interface was plain, while the other included a happy face. Regardless of concerns, users were more likely to disclose when experiencing a positive affect. As aptly summarised by the researchers, “*rationality in privacy-related decision-making is bounded by psychological limitations*” [202]. We now move forward to consider other factors for the Paradox.

Design and control

Constraints. When individuals interact with systems, they usually do so through user interfaces. Furthermore, to use protective tools, additional menus must be navigated. Therefore, the design of these portals can have an influence on privacy behaviour [59]. Often the interests of the user and the vendor might not be aligned. Whereas the former might wish to protect their data, the latter would prefer increased disclosure [334]. Therefore, to limit the influence of privacy, certain strategies can be implemented. Bösch et al. [59] summarised a range of ‘dark patterns’. For example, ‘bad defaults’ are frequently included on modern technologies. Since users rarely adjust their default settings [244], data can be extracted through lax configurations.

Control. There is also ‘privacy zuckering’, where options are extensive but opaque. Vendors can then claim to provide protective tools, even though they are too complex to use [59]. This also helps to grant consumers control, which can have counter-intuitive results. Brandimarte et al. [62] studied the influence of control over disclosure behaviour. Surprisingly, when participants were given additional settings, they revealed more information. This could be related to risk homoeostasis, where individuals become daring when given protection [392]. Kearney and Kruger [201] studied homoeostasis in the context of information security. When users felt protected from hackers, they often compensated with riskier activities. Therefore, even if individuals see privacy settings, they might engage in frequent disclosure.

Resignation. Users can become resigned if they doubt the efficacy of protection. Marwick and Hargittai [254] conducted 10 focus groups with a total of 40 students. These discussions explored the incentives and disincentives to online sharing. Since

many participants felt that violations were inevitable, they sought little protection. The authors conducted additional focus groups, exploring privacy awareness and protective ability [166]. Users expressed that settings were temperamental, complex and liable to change. Even when they did make use of protection, they doubted its efficacy. This led to privacy settings being frequently neglected. Even if individuals have concerns, they might not act unless it is thought feasible.

Social norms

Concern. As previously mentioned, there is a western social norm for privacy [352]. Huang and Bashir [185] found that most people consider it a human right, regardless of their culture. Through social desirability bias, individuals are encouraged to express opinions that conform with norms [138]. This might result in inflated claims of concern, hence presenting the Paradox. To reduce this risk, our studies implement techniques to minimise the bias. For example, since concerns can increase when privacy is primed [65], we frequently include decoy questions. Furthermore, we seek to disguise the topic to mitigate demand characteristics [287]. These occur when participants adjust their responses to match the study purpose.

Behaviour. However, norms do not merely affect privacy concern. There is also an expectation for individuals to use modern technologies. Many platforms, particularly social networking sites, are considered important for social capital [127]. As individuals begin to use portals frequently, data disclosure becomes a normal routine [110]. Gratification is then received from accessing such services, even if privacy is placed at risk [313]. People might initially be reluctant to reveal large quantities of information. However, Snyman et al. [351] found that users are prone to the ‘lemming effect’. They will follow dangerous group norms, even if these conflict with their personal concerns. Finally, the framing of the topic can also be influential. Politicians claim that citizens should have ‘nothing to hide’ if they do no wrong [353]. If individuals are pressured to both claim concern and avoid protection, it is little surprise the Paradox exists.

Privacy awareness

Knowledge. While no user is ‘average’, most individuals do not possess a technical education. Therefore, they cannot be expected to have expertise in security or privacy. Indeed, many studies have demonstrated a lack of awareness. We define ‘privacy awareness’ within our Personal interpretation subsection. Bashir et al. [49]

distributed a two-part questionnaire to 455 individuals. Participants were found to have little awareness of many privacy threats. Furthermore, few expressed knowledge of protective settings. However, they still continued to use a range of online services. Madejski et al. [247] compared the intentions and privacy settings of 65 participants. Each user made at least one error, demonstrating the commonality of misunderstandings. Finally, a 2017 UK survey found that only 10% claimed to know privacy practices [99]. In their poll of 2,153 adults, they also discovered 8% understood data sharing. If individuals lack knowledge of risks and protection, they are unlikely to guard their data [48]. As outlined earlier, there is no binary distinction between awareness and ignorance. However, many users tend to have a low degree of privacy understanding.

Absence. It is natural that users lack this awareness. Privacy is often considered a secondary goal [188], and other topics may be more important to the individual. Because of this, users might display ‘rational ignorance’ in lacking knowledge [121]. This doctrine compares a risk against the effort required to become informed. If a threat is deemed to be low, individuals will not invest time in privacy policies [397]. However, without a degree of understanding, risks cannot be perceived accurately [350]. This might lead individuals to underestimate the importance of privacy. Furthermore, when understanding is low, the benefits of disclosure tend to be overestimated [157]. If knowledge can be provided, users should be more motivated to use protection.

Influence. The Paradox is often found when awareness is lacking [48]. This is sometimes due to ‘privacy cynicism’, where individuals doubt the efficacy of protection. Hoffmann et al. [179] explored the matter through focus groups with 96 participants. Four recurring themes were identified: uncertainty, powerlessness, mistrust and resignation. To avoid cognitive dissonance, users can distrust settings when they do not understand them. Oetzel and Gonja [284] considered the Paradox through social representations. They highlighted that while offline privacy is common knowledge, individuals lack an understanding of digital value and risk. Baruh et al. [48] conducted a meta-analysis of awareness and Paradox factors. They found that concerns were not a strong predictor of user behaviour. However, when users had ‘privacy literacy’, they were more likely to use protection. This demonstrates promise for reducing the Paradox.

schraefel et al. [335] deemed awareness, among other factors, to affect the Paradox. They described how individuals are “*sense-making*” and will assume that data is required if requested. Users are too busy to investigate each request, and the question delays their primary goal. Therefore, they assume the best and agree to oft-opaque

terms. However, these individuals are seldom aware of how their data is used. When the true nature of the transaction is highlighted, the abstract issue becomes concrete. Users then frequently express concern, creating this Privacy Paradox [335]. If we can inform individuals in advance, they might limit their data disclosure. This could reduce a disparity between concern and behaviour.

Approaches. Previous work has suggested the feasibility of this technique. Pötzsch [309] developed requirements for awareness-enhancing systems. It was essential these tools were understandable, personalised and performant. They believed that education could reduce unintentional actions. Therefore, since individuals should possess concerns, this would address the Privacy Paradox. Deuker [114] came to a similar conclusion in their 2009 work. They discussed how incomplete information could contribute to unexpected violations. To address both this and cognitive issues, additional details should be provided. Deuker also suggested that education target two components: privacy risks and protective techniques. Although these studies did not collect data, we use similar rationale in our own experiments. We believe this sensible, since systematic reviews recommend the enhancement of awareness [153].

It should be noted that while aforementioned works have studied the Paradox, few have attempted mitigation. Jackson and Wang [194], to be outlined later, reduced the disparity within a session. However, no research has mitigated the issue over a span of time. As we appear to achieve this in a two-month study, our work presents a novel contribution.

Criticisms and scepticism

Attitude-behaviour gap. While there is much evidence for the Paradox, it has been criticised on a number of grounds. Some claim that the concern-behaviour disparity is not a remarkable discovery [14]. The attitude-behaviour gap is commonly found in psychology, with intentions often diverging from actions [14]. Ajzen and Fishbein [14] considered the relationship between attitudinal predictors and individual behaviour. They found there was often little correspondence between the two components. On these occasions, significant associations were rare. There is also the notion of the value-action gap [45]. Peoples' values can be dynamic and contradictory, with behaviour restricted by external constraints. If concerns and actions often misalign, why is privacy interesting?

In response, it is true that attitudes and behaviour are frequently misaligned. However, this is often due to methodological issues. Our studies adopt the *principle of compatibility* [12], where queries correspond to the same topic. This is achieved

through contextualised questions and defined actions. In cases of compatibility, Ajzen and Fishbein [14] found that significant relationships were common. Since we still confirmed the Paradox, the issue is more unorthodox. Furthermore, while hypocritical behaviour is frequent, privacy presents uncommon challenges. Firstly, few users are aware of their personal risk [87]. Secondly, even if users desire protection, knowledge is relatively rare [49]. Thirdly, since privacy is intangible [195], improvements might not be noticed. Fourthly, the interests of users and service providers can be in tension [334]. Finally, social norms push for both greater concern [352] and greater disclosure [127]. If the Paradox can be mitigated, it presents a rare success.

Instinctive concern. As previously mentioned, individuals might intuitively oppose privacy violations. However, once they consider their opinion, their response might adjust [303]. This suggests that concerns might be driven by social expectations or ‘gut feelings’. Further research claimed to dismiss the Privacy Paradox. Baek [39] used counterarguments to debate with their participants. Through three studies, they found that concerns were highly superficial. If individuals had time to consider the matter, perhaps they would express different views.

This appears to call into question the usage of questionnaires. However, such instruments are a standard approach to gauge opinion. They are also a trusted and popular technique. Therefore, aforementioned results may be considered an idiosyncrasy of the Paradox, rather than a methodological fault. If most participants initially respond in the same manner, one would surmise that their view is prevalent. Furthermore, when opinions are challenged or debated, it is natural for them to adjust. Despite these points, we do implement measures to reduce the risk. Most of our forms request qualitative rationale, prompting users to consider their justifications. We also make use of semi-structured interviews (Chapters 6 and 8), where opinions are determined from discussion. Through such conversations, participants can express their considered concerns.

Concern concepts. Some have criticised that ‘privacy’ is considered as a single concept. Dienlin and Trepte [116] expressed that it should be subdivided into ‘informational concerns’, ‘social concerns’ and ‘psychological concerns’. When they compared holistic opinions against behaviour, the Paradox was confirmed. However, when the elements were analysed individually, the issue vanished. In our studies, we explicitly scope our focus to ‘information privacy’ [92]. Both concerns and behaviour relate to the usage of technological devices. Therefore, when confirming the Paradox, it should not be due to this issue.

New Paradox? Adorjan and Ricciardelli [9] asserted that a new Privacy Paradox was emerging. In their focus groups with teenagers, they found that many claimed to have ‘nothing to hide’. Despite this, the children frequently used Snapchat to achieve privacy. In this manner, their participants used greater protection than their concerns suggested. While we recognise the rationale, we doubt that the Paradox will reverse. Young people are more open with their data than older individuals [89]. Furthermore, research suggests that ‘nothing to hide’ is a minority viewpoint [185].

Sampling biases. As expressed by Preibusch [310], the Paradox might arise if different populations are inappropriately compared. One poll might suggest that 90% of respondents are concerned about privacy. A separate survey could indicate that only 20% adjust their app permissions. While these metrics can imply a trend, direct comparisons are improper if samples differ. To address this, we analyse the concerns and behaviour of the same participants. Furthermore, the Paradox is only ever studied on an individual level. This allows opinions and actions to be compared in a defined context. Therefore, when we assess Paradox prevalence, our analyses are specific to each respondent.

Context. Since privacy is contextual [278], concerns might differ based on the topic. Similarly, users might use more protection in certain circumstances. Trepte et al. [369] criticised previous work for comparing abstract opinions with practical behaviour. For rich analyses of the Paradox, both components should relate to the same context. Therefore, we take measures to ensure our studies are adequately grounded. After the Paradox has been confirmed (Chapter 4), we deconstruct the issue through semi-structured interviews. Concerns are gauged through customised questions, while behaviour relates to the same context. This approach is then also used in Chapters 7 and 8. Furthermore, participants own the concerned devices in the final three studies. Therefore, they should be able to provide responses in the appropriate context.

Personal interpretation

As discussed in the above subsections, there are many factors which contribute to the Paradox. However, in our thesis, the topic is assessed through the following lens. We believe that most people aspire to privacy as an abstract principle [122]. While some might reject the concept, it is strongly encouraged by western social norms [352]. Although it is impossible to have perfect protection, individuals would like to be as safe as possible. However, their efforts are impeded in a number of ways. Firstly, platforms are developed to discourage protection [59]. This can contribute

to apathy, as individuals might regard privacy as too much effort. Secondly, social norms support the disclosure of personal information [127]. Thirdly, and we assert most importantly, individuals often lack privacy awareness [49].

While awareness might appear a simple topic, it has great complexity. In Pötzsch's Paradox work [309], she defined it as "*the attention, perception and cognition of:*

1. *Whether others receive or have received personal information about him/her, his/her presence and activities,*
2. *Which personal information others receive or have received in detail,*
3. *How these pieces of information are or may be processed and used, and*
4. *What amount of information about the presence and activities of others might reach and/or interrupt the individual."*

We agree that the term can refer to a range of components. A person might possess or lack awareness of a privacy threat. For example, a Facebook user might not realise that their settings are permissive. This differs from an awareness of protective features. Even if an individual recognises their vulnerability, they might not be able to act. Similarly, they might know of privacy settings but assume that their data is safe. This complexity supports our usage of Protection Motivation Theory (PMT) [328], as described in Section 2.5. In this model, awareness of the threat is distinct from awareness of the response. Through our educational games, we seek to support both components.

However, we consider awareness to be distinct from education. A person might know of an issue without understanding the matter. For example, a smartwatch user might recognise that their location is monitored. But they may not know why it occurs or how to prevent it. We also distinguish awareness from skill. In the latter case, an individual possesses "*a particular ability*"¹. Although Facebook users might understand protection, they could experience difficulty navigating their settings. A skilled individual would have greater confidence in configuring the interface. Considering Pötzsch's definition and the contrast with education and skill, we define 'privacy awareness' as the following: "*the state of being aware of a privacy threat/vulnerability and/or the existence of protection*". This builds upon the Merriam-Webster dictionary definition². It is by increasing awareness, and providing information and practice, that we seek to encourage protective behaviour.

¹<https://en.oxforddictionaries.com/definition/skill>

²<https://www.merriam-webster.com/dictionary/awareness>

We do not wish to influence or coerce individuals. If a person does not desire protection, then they should not implement it. However, it is important that users can make informed decisions. When risks and settings are not understood, then actions might be unwise. By offering education and practice, they have an opportunity to protect themselves. Although awareness is often touted as a solution [114,309], the Paradox cannot be wholly eliminated. Individuals might prefer the benefits of disclosure to the advantages of protection [107]. Even if privacy is genuinely appreciated, cognitive biases will impede ‘perfect’ behaviour [6]. Therefore, our work does not aim to eliminate the Paradox. In contrast, it seeks to support informed decisions, which should result in Paradox mitigation.

2.3 The Internet-of-Things (IoT)

Overview. The Internet-of-Things (IoT) has been described as the digital revolution of the twenty-first century [377]. It promises to connect a vast number of ubiquitous components, enmeshing itself in our daily lives. As such it offers a wealth of opportunities for productivity and convenience, and is predicted to generate trillions of dollars for the global economy [251]. It has grown enormously over the past seven years, with connected devices quadrupling in this decade [129]. 55 billion products are predicted by 2025, indicating the influence the IoT will have on our societies [276].

However, with the concept rapidly-becoming clichéd, we should remain grounded with a clear definition. In their 2005 annual Internet Report, the ITU [193] described the network as “*anytime, anywhere, by anyone and anything*”. While this phrase encapsulates the vision behind IoT developments, it lacks a degree of precision. Miorandi et al. [268] define it as a “*global network interconnecting smart objects by means of extended Internet technologies*”.

While this definition is appropriate, all descriptions are challenged by the IoT’s heterogeneity. It comprises a nebulous collection of technologies, ranging from wearables to smart appliances to home automation systems [215]. Due to this breadth, scoping was essential to our research. When the topic is first considered in Chapter 5, we constrain our focus to consumer products. These devices are both recognisable to the public and amenable for user studies. As our analyses progress, we then focus directly on smartwatches. While they are discussed in the next section, their selection is justified in Chapter 7.

Advantages. As can be expected from such developments, the IoT offers many benefits. Consumers might reduce their electricity consumption through smart me-

ters. Homes could monitor the safety of the elderly, alerting families if an issue occurs. Factories could increase their output, while citizens navigate the roads in self-driving vehicles. However, despite these possibilities, the IoT presents real risks to privacy.

Privacy issues

Data collection. The IoT promises to connect together vast numbers of computing devices. These products will interact both with each other and their surrounding environments. Therefore, it is likely that large quantities of data will be collected. As expressed by Perera et al. [300], “*the amount of user data ... will be significantly higher than in the past*”. Since many products lack in-built capabilities, this information might be transmitted remotely for processing [300]. Whereas one’s data used to reside in their hard drive, now details are stored in an expanding cloud. Hölbl et al. [181] presented an overview of the IoT’s security and privacy challenges. Chief amongst these issues was the collection of sensitive data. Aktypi et al. [15] constructed a risk exposure tool for the IoT. They found that personal details, such as home address, could be inferred from wearable data.

Data is used to provide IoT functionality [301]. Smartwatches might monitor your location to deliver navigational services. Home automation systems might adjust settings in response to environmental sensors. While this can be beneficial, it further demonstrates the frequency of collection. Data access might also be encouraged by financial incentives. Information could be lucrative and sold to advertising networks [334]. This can subsidise the cost of the products, hence supporting market expansion.

Ubiquitous monitoring. Although the online-offline divide is increasingly blurred, traditional monitoring has been largely ‘virtual’. When browsing the web, platforms can track identities through cookies or IP addresses. However, since many IoT devices possess sensors, they may access additional details. Your smartwatch might know your location, while homes can monitor their indoor environment [301]. Elkhodr et al. [126] conducted a literature review, considering the risks posed to location privacy. They expressed that “*it is almost impossible to achieve perfect privacy as long as seamless communication is taking place*”.

If device monitoring is recognised, then users have an opportunity to limit it. Unfortunately, surveillance is often invisible and surreptitious. Smart TVs have been found to overhear conversations, before sending transcripts to remote locations [154]. The Amazon Echo has also raised concerns that our environments are being monitored [273]. Ford and Palmer [142] analysed Alexa network traffic over a 21-day period.

They found that conversations can be inadvertently processed by these services. If such gadgets become normalised, privacy might be reinterpreted.

Resource constraints. Larger IoT technologies, such as home automation systems, might possess computing power. However, many smart devices are constrained in the operations they can perform. Due to the size and mobility of smartwatches, they may be particularly affected. In small form factors, there is little space for extensive storage. Therefore, data might be transmitted for remote processing [300]. Furthermore, when batteries are limited, computational complexity is constrained [372]. This poses particular risks to confidentiality, as cryptographic algorithms require complex operations [372]. Many IoT protocols lack standardisation, with a variety of techniques proposed [261]. This can further threaten the privacy of data.

Privacy Paradox considerations

As discussed, the IoT could pose several issues for privacy. When exploring the Paradox, we seek to assess individuals' concerns and actions. For several reasons, we assert that IoT will contribute to lax behaviour. This could exacerbate the disparity between the two components.

Functionality. We believe that consumers might be preoccupied by functionality. Rheingans et al. [319] found that activity trackers are considered to be hedonic technologies. In this manner, decision-making is driven by pleasure rather than other factors. Since the IoT promises convenience, we expect this environment to be judged similarly. In her 2016 publication, Bailey [41] agreed with this conclusion. She opined that users were “*seduced by technology*” and unaware of company practices. Even Brill, an FTC commissioner, expressed that concerns would not prevent IoT adoption [67].

As previously outlined, hyperbolic discounting can influence the consideration of risk and reward [5]. In IoT environments, we expect individuals to favour short-term gratification over long-term protection. This was supported by the 2017 literature review by McCloud et al. [258]. They considered the influence of wearable devices, and concluded that emotion could bias privacy decisions. Aleisa and Renaud [21] came to a similar judgement based on empirical evidence. They explored IoT awareness and concerns through a 236-person survey. It was found that participants had little knowledge of what their devices collected. Even when they possessed concerns, convenience appeared to outweigh this factor. Based on these issues, we expect IoT owners to use little protection. If concerns remain, this might contribute to the Paradox.

Usability and unfamiliarity. We believe that protection will be constrained by poor usability. While the IoT is expanding rapidly, it is still a novel environment.

Furthermore, due to its heterogeneity, devices are produced by a wide range of vendors [268]. This will contribute to a miscellany of manufacturers and user interfaces. Since mental models might not be aligned to new interactions, errors may occur [317]. This could be especially true due to the unfamiliarity of smart devices [177]. Several studies have criticised the usability of the IoT. Foster et al. [143] conducted a report into smart energy integration. They found that many interfaces were too complex, and these posed a barrier to adoption. Smartwatches also appear to be limited in their usability. Chun et al. [90] studied device usage over a one-week period. Through their surveys and interviews, they concluded that rich interactions are constrained. Privacy settings tend to be open by default [59], and users rarely adjust these options [244]. We believe this will impede protection, further contributing to the Paradox.

Privacy Paradox studies

Previous work. We now consider relevant work in this area. To the best of our knowledge, the Paradox has not been studied within the ‘Internet-of-Things’. However, several publications have considered the relationship between the phenomenon and the environment. Maple [252] conducted a literature review of IoT privacy. In the document, he discussed how changing social norms will challenge protection. He also highlighted that, through smart devices, data collection might be greater than in the past. Although an impact on the Paradox was mentioned, empirical work was not undertaken. Barth and de Jong [47] considered the topic when undertaking their 2017 literature review. They explained how mobile computing can require faster decisions, thus placing pressure on the Paradox. However, the phenomenon was not studied in this new environment. Andrushevich et al. [25] discussed the integration of IoT systems in the BUTLER project. In their paper, they outlined the influence of increased (data) sensitivity and availability. Although they used the term ‘Privacy Paradox’, they actually described the ‘Personalisation-Privacy Paradox’ [10]. This is a different topic, which balances service customisation against user concerns.

While the Paradox has not been directly studied, privacy calculus has been adopted for some analyses. Kowatsch and Maass [218] examined four IoT domains: transport payment, navigation, smart home and health monitoring. They used the Extended Privacy Calculus Model to predict usage intention. Since no negative relationship was found between risk and intent, this might suggest a disparity. Derikx et al. [111] explored the influence of concerns on connected car services. They surveyed 55 participants through discrete-choice questionnaires. Although respondents preferred non-invasive systems, they would disclose data for a financial incentive. Lee

et al. [233] used privacy calculus in their 2018 publication. Several IoT technologies were considered, including connected appliances and smart grids. Through their 300-person survey, they found concerns actually increased usage intention. Although the Paradox was not discussed, a disparity might exist between opinions and behaviour.

Scarcity. As demonstrated, few studies have considered the Paradox in the IoT. We believe this is likely for several reasons. Firstly, the issue is still being explored in less-novel contexts. Publications continue to focus on social networks [161] and mobile phones [194]. This is quite understandable, since these interfaces are more amenable to analysis. The Paradox remains far from explained, with additional interpretations arising each year [8, 161]. Until the matter is thoroughly explored, it is unlikely to venture into the IoT. Secondly, the environment is still relatively novel. Although smart devices are increasing in prevalence, they are less common than traditional computers. As the market further expands, IoT studies will become more popular. Finally, the Internet-of-Things is challenging to analyse. It poses issues in terms of heterogeneity [215], usability [143] and familiarity [42]. While Facebook can be assessed through browser extensions [384], IoT analyses are more constrained.

Importance. Despite this research gap, it is important that the Paradox is explored. As discussed, IoT privacy issues are well-documented. Due to the functionality of devices, user data is frequently collected. If consumers cannot protect themselves, they might be placed at risk. The IoT is also expanding rapidly, with 55 billion products predicted by 2025 [276]. When these technologies pervade our environments, privacy might be threatened. Finally, since the IoT is novel, users are even less likely to possess privacy awareness. This might create a greater disparity. To support informed decisions, we explore this nascent environment.

2.4 Smartwatches

Selection. We begin the thesis by discussing the Paradox in the context of the IoT. However, as our studies progressed, we were necessitated to scope our focus. The environment is too heterogeneous to develop broad mitigative approaches. Smartwatches were selected for several reasons, as detailed in Chapter 6. Of the IoT products we analysed, these devices appeared most prone to the Paradox. Since we wished to support privacy, we turned our attention to this context.

Overview. Smartwatches are defined as “*an electronic wristwatch that is able to perform many of the functions of a smartphone*” [97]. Initial concepts date back

several decades to the 1980s³. These Ubicomp systems sought to provide a user with wearable functionality. However, wrist-based computing only became popular in the early 2010s³. Exercise trackers, such as the Nike Fuelband, were first released at this point. Although these are indeed ‘wearables’, we distinguish such products from true smartwatches. Trackers seek to monitor exercise and tend to lack most smartphone functions [381]. Smartwatches have advanced operating systems and often support third-party apps. Since this distinction is commonly made [208,381], we believe our scoping to be appropriate.

Advantages. Smartwatches clearly provide a wide range of benefits. Since they reside on the wrist, they offer a quick shortcut over phones. Through the use of brief notifications, an individual can save time. They also often provide exercise tracking and heart monitoring. Furthermore, newer models support the installation of third-party apps. This further extends the functionality of the device. Smartwatches are growing in popularity, with the market expected to blossom over the next five years [226]. However, they are not without privacy risks. We now discuss the issues that this environment presents.

Privacy issues

Sensitive data. To provide such convenient services, watches require access to sensitive data. These devices can be defined as either paired or stand-alone. In the former case, details are synchronised between the wearable and a nearby smartphone. In the latter case, the watch acts autonomously from other products. Sensitive data can be stored and accessed in either situation. Details range from text messages to phone contacts to calendar appointments [118]. The watch could also contain GPS readings and heart rate records [16]. If privacy permissions are not restricted, apps can gain access to this information. Users are likely to avoid these permissions as they can limit their device’s functionality [335]. Furthermore, product loss/theft could place these details into the hands of strangers. Therefore, it is important that individuals use their privacy settings.

Data collection. Smartwatches can also collect data from their surrounding environment [125]. Virtually all models have GPS access, whether natively or through a paired phone. In many cases, the location service is enabled by default. While this offers great advantages to functionality, it allows positions to be identified in real time. Heart-rate tracking is also provided by a range of smartwatches. This is useful

³<https://www.wearable.com/smartwatches/smartwatch-timeline-history-watches>

for those who wish to gauge their fitness. However, such metrics are deemed to be ‘Protected Health Information’ under HIPAA [373]. The monitoring might not even be noticed, since watch settings are rarely checked [371]. In such cases, users will have little knowledge of their privacy.

Device security. The quality of smartwatches has improved over the past decade. However, these devices often remain prone to attack. For example, HP demonstrated that most models possessed major vulnerabilities [176]. They analysed 10 leading watches and found communications could be intercepted in 90% of cases. Do et al. [118] conducted a study on Wear OS⁴ products. They proposed a technique to compromise a device’s boot chain. Through an empirical analysis, they extracted sensitive locations, text messages, health details and voice recordings. These files could both reveal personal data and support further inferences [106]. Baggili et al. [40] undertook forensic analysis of Samsung and LG watches. They first forced root access, before downloading data from the devices. They succeeded in extracting events, contacts, emails and text messages. Although smartwatch security has developed, data might still be vulnerable to an adversary.

Lack of protection. Previous work suggests that users express concern for wearable privacy [162,319]. Since smartwatches contain sensitive details, one might expect protection to be adopted. However, we believe that privacy settings will be rarely used. This is for several reasons. Firstly, such options are rarely configured in other environments. This holds true for both social networking sites [242] and mobile phones [22]. Secondly, smartwatch configurations appear to be particularly challenging. Horcher [183] studied the security usability of Apple Watches. She found that all participants struggled to enter a PIN. Finally, research has demonstrated that protection is rarely used. Udoh and Alkharashi [371] explored the concerns and behaviour of 10 students. Through a survey, they found most took no action to protect their smartwatches. Based on the above points, we expect this to be true of many users. Since protection is rarely used, the Paradox might be prevalent. However, previous work has not analysed empirical privacy behaviour.

Privacy Paradox studies

Previous work. The Paradox is also rarely considered in this environment. A notable exception is the analysis by Hallam and Zanella [162], which makes use of Construal Level Theory. They surveyed 103 individuals to explore the roles of perception

⁴‘Wear OS’ is formally known as ‘Wear OS by Google’. It was recently rebranded from ‘Android Wear’, and older papers might refer to this name.

and behaviour. While short-term benefits were inflated, long-term concerns appeared discounted. Privacy awareness also reduced disclosure, suggesting our proposed approach might be effective. While this work highlights the Paradox, it is quantitative and focused on healthcare wearables. As in the case of the IoT, privacy calculus is studied more frequently. Li et al. [237] developed a model which was validated by 333 device owners. They also considered healthcare products, but included fitness trackers in this definition. It was found that rational risk-reward calculations were made, conflicting with the Paradox. However, functionality might be more gratifying on consumer smartwatches.

As previously mentioned, Rheingans et al. [319] found activity trackers are driven by considerations of pleasure. The researchers assessed risks and intention through a 115-person survey. Since concerns appeared to have little impact, they claimed that vendors should highlight functionality. This further suggests that the smartwatch market is driven by convenience. Wieneke et al. [390] explored why wearables have popularity despite the risk. They conducted interviews with 22 users, finding low awareness of the threats. While they used privacy calculus, they concluded that the decision-making was not rational. This might partially explain the Paradox, although smartwatches were not specifically considered. These studies all analysed the adoption of technology, rather than wearable behaviour. Our research differs, by both exploring and encouraging protective actions.

Scarcity. In a similar manner to the IoT, smartwatches appear rarely studied. We believe this is predominantly due to three reasons. As before, attention continues to focus on less-novel products. While the issue is being analysed on mobile devices, there is less reason to target new areas. Secondly, smartwatches are novel and still lacking in popularity. Although their market has grown [191], they are not as common as many technologies. Once security/privacy behaviour is studied, Paradox analyses will soon follow. Finally, smartwatches are constrained in several regards. They possess small screens, few buttons and limited processors. Their interfaces also tend to be proprietary, impeding the collection of data. Therefore, other environments are more amenable to study.

Importance. However, we believe it is important that the Paradox is explored on smartwatches. The devices can house sensitive details, ranging from text messages to calendar appointments [118]. This data could be accessed by app companies if permissions are not adjusted. Unfortunately, it appears as if protection is rarely used [371]. This is consistent with other devices, since permissions are abstract and opaque [335]. To support informed decisions, we seek to increase privacy awareness.

Mobile computing has been claimed to encourage irrational decision-making [47]. Since interactions are fast and fluid, users might forget about privacy protection. Furthermore, we expect smartwatches to grow in complexity. As their capabilities and popularity increases, data collection will likely escalate. If we wish to promote protection, analyses must be conducted before it is too late.

2.5 Behaviour Change

Ideology. Our research seeks to mitigate the prevalence of the Paradox. To achieve this, we wish to encourage the greater usage of privacy protection. It is true that behavioural adjustments do not solely affect the issue. If concerns also increase, the disparity between the factors might remain the same. However, since opinions are often based on principle [278], we do not expect them to radically change. Furthermore, when users begin to understand protection, their unease might decrease. Therefore, our research seeks to encourage protection rather than reduce concerns.

To achieve this, we aim to persuade our study participants. This explicitly differs from influencing or coercing the individuals. Persuasion is defined as an “*attempt to change attitudes or behaviours or both (without using coercion or deception)*” [141]. By highlighting the threats to privacy, we hope to increase awareness. Since awareness also concerns protection, we seek to present protective features. If individuals do not wish to use settings, that is their decision. Even when risks are salient, a person might be drawn to functionality. We reject paternalism, as a user will understand their situation better than a researcher. However, since we find awareness to be the greatest issue, information should support protective behaviour.

Behavioural theories

Introduction. Behaviour has been studied intricately through the field of psychology. As the discipline has developed, several theories of action have become influential. Since they seek to explain decision-making, they are of importance to our work. We discuss popular models before selecting the theory most apt for our research.

Theory of Reasoned Action (TRA). This theory seeks to predict behaviour based on existing attitudes and intentions [136]. An illustration can be found below in Figure 2.1. It first explores ‘behavioural beliefs’, considering opinions about a potential action. For example, a person might believe protection is provided by app permissions. It then examines ‘outcome evaluations’, in which individuals consider whether the result is beneficial. In the above case, a user might think that protection

will reduce their vulnerability. These factors are considered alongside the ‘subjective norms’. The norms are themselves informed by ‘normative beliefs’ and ‘motivation to comply’. The former concerns the societal expectations surrounding an action. For example, an individual might sense a norm to value their privacy. The latter relates to the willingness to abide by that norm. In our example, the person might wish to conform with expectations. As a result, they should intend to protect their data.

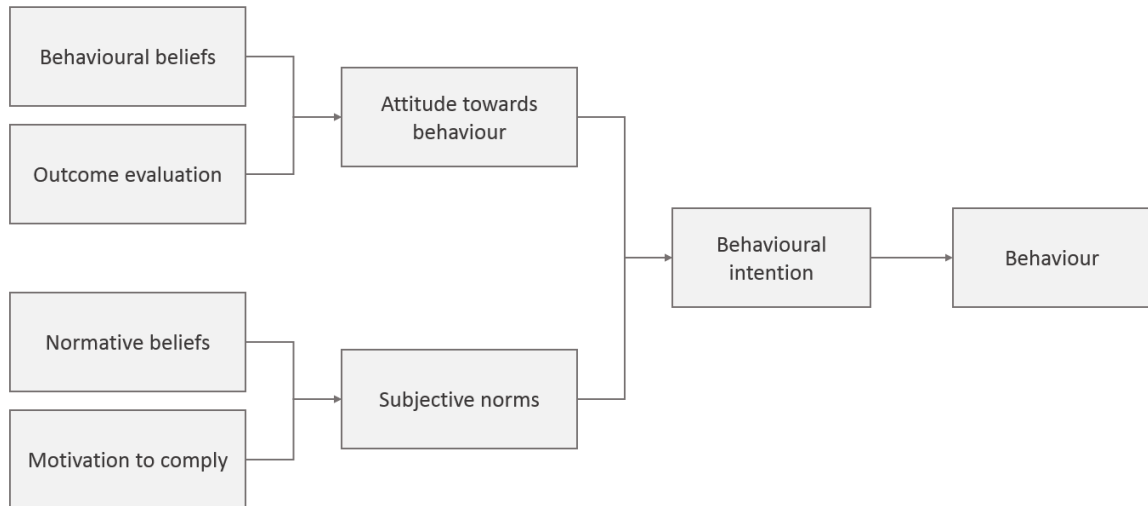


Figure 2.1: Theory of Reasoned Action (TRA) [136]

The theory has been used in previous security and privacy research. Siponen [349] applied the model, alongside other theories, to organisational security awareness. He developed a persuasion framework to enhance employee compliance. However, it was not evaluated empirically. In follow-up research, Siponen et al. [348] conducted a survey with 917 IT professionals. Results were analysed through Structured Equation Modelling (SEM), before being applied to several theories. An intention to observe policies was found to have a significant impact on compliance.

This theory does not appear appropriate for our research. Individuals might value their privacy and wish to comply with the norms. However, they may still fail to take action. This is partially because the model does not consider behavioural constraints. It assumes that all behaviour can be controlled, which is rarely the case [66]. Self-efficacy is omitted: the confidence a person has in their own ability [43]. If users are not confident they can update their settings, they might choose to avoid them. Furthermore, TRA excludes the consideration of behaviour requiring skills [240]. Since device protection is not trivial, the theory might lack applicability.

Theory of Planned Behaviour (TPB). This theory [13] is illustrated below in Figure 2.2. It builds on TRA by including attitudes and subjective norms. However,

it differs by introducing ‘perceived behavioural control’. By considering this factor, the model is more appropriate for our privacy studies. For example, a person might value the advantages of guarding their data. Protection might also be encouraged by social norms [352]. However, if they lack self-efficacy, they might fail to act.

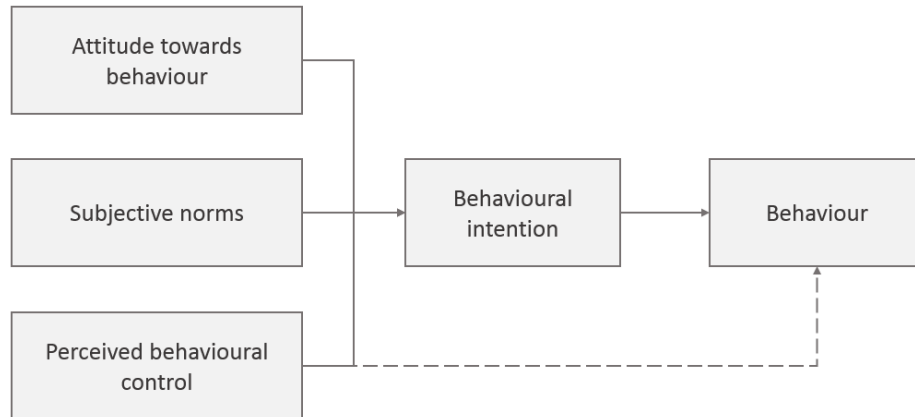


Figure 2.2: Theory of Planned Behaviour (TPB) [13]

The theory has been applied in privacy research, exploring how salience influences data disclosure [189]. 45 participants took part in the study, with these individuals split between 5 groups. Behavioural attitudes were influenced by colour-coding, while norms were targeted through peer advice. The approaches appeared successful, since the treatment groups disclosed less personal information. Yeo et al. [395] used the theory to raise awareness of security protection. Their participants completed a TPB questionnaire in pretest and posttest. After using a persuasive technology, their protective attitudes were seen to increase.

However, the model does possess deficiencies. Risk is central to our research, as it influences both concerns and behaviour. Although the theory considers outcomes, the concept of risk is not salient. As outlined by Norman and Conner [281], it fails to account for susceptibility or severity. A person be confident in guarding their data. But they may lack motivation if the threat is minimal. Furthermore, while self-efficacy is explored, the efficacy of the action is not considered. Even if a person is skilled, they are vulnerable if responses are ineffective.

Protection Motivation Theory (PMT). This model follows a different approach [328]. Its high-level structure is illustrated below in Figure 2.3. At a high level, the ‘threat appraisal’ is balanced against the ‘coping appraisal’. If the risk is great and the remediation is simple, individuals should take action. The threat is composed of perceived vulnerability and perceived severity. If a person feels susceptible to a serious issue, they might guard their data. This threat is subtracted from

rewards, which encapsulate the benefits of avoiding protection. For example, if apps offer convenience then users might neglect their permissions.

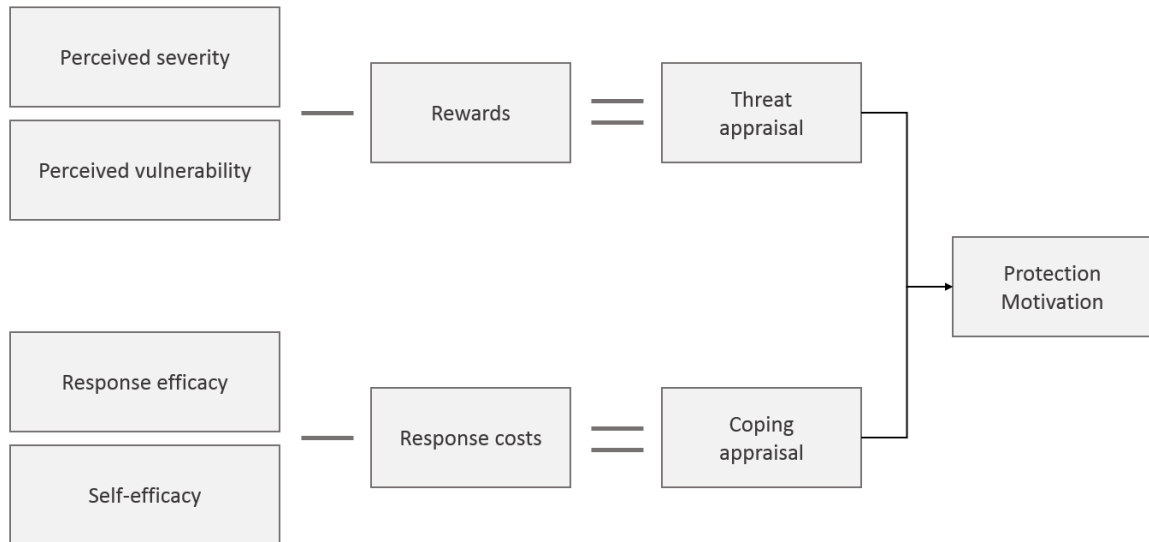


Figure 2.3: Protection Motivation Theory (PMT) [328]

Efficacy comprises both the response efficacy and self-efficacy components. The former applies to the action’s success while the latter concerns the user’s abilities. A person might be confident in charging their watch, but this would not protect their privacy. These factors are balanced against response costs. They relate to how much effort the action takes to complete. An individual may understand passwords and appreciate their benefits. However, if they must be typed on a small screen, protection might be deterred.

This theory has proved successful in privacy research. Chenoweth et al. [86] used the model to predict the usage of protective technologies. They studied how the factors would affect the intention to use anti-spyware tools. Through an analysis of 232 participants, perceived vulnerability and response costs appeared to be influential. Another study used PMT to assess how concerns impact protective behaviour [396]. This was undertaken through a questionnaire distributed to 144 students. Perceived vulnerability was most important, though perceived rewards also had significance. Although these works did not explore the IoT, they align well with our research.

This model appeared most appropriate for several reasons. Firstly, it considers concerns and responses, similar to the structure of the Paradox. Through this, we can assess why users fail to protect themselves. Secondly, since it encompasses both risk and protection, it has relevance to our privacy research. Thirdly, it has been recommended for behaviour change in cyber security [66]. Finally, it has been deemed

highly for evaluating behaviour. Floyd et al. [140] conducted a meta-analysis of 30,000 participants over 65 studies. They found that the PMT components had a good effect size. Since this theory appears most appropriate, it was adopted in our research.

Privacy by Design

Studies have long sought to encourage privacy-protective behaviour. In our research, we make use of educational smartwatch games. We first introduce the alternatives, their approach, and why they were not selected.

Overview. Privacy by Design (PbD) seeks to encourage the inclusion of privacy into initial plans [82]. Through this, protection is designed and implemented, rather than overlooked. It also ensures that privacy is considered at the start, rather than retrofitted into existing systems.

A number of frameworks have been developed based on the concept. PbD was originally formulated by Ann Cavoukian, following a Canadian-Dutch collaboration in 1995. It consists of seven principles, including ‘privacy as the default’, ‘full functionality’, and ‘privacy embedded into design’ [83]. For example, users should not have to opt-in to receive data protection. Langheinrich [227] also proposed six principles for PbD. These included ‘choice/consent’, ‘anonymity/pseudonymity’ and ‘access/recourse’. They were suggested for ubiquitous systems, but have found implementation in a range of fields. Frameworks have even been developed for IoT environments. Perera et al. [299] compiled 30 guidelines, adapting the eight design strategies of Hoepman [178]. They recommended that data acquisition, storage and retention be minimised.

Issues. However, PbD is not a panacea. Firstly, the principles are directly in tension with corporate interests [334]. As Spiekermann aptly expresses, privacy can “*limit strategic options and impact a firm’s bottom line*” [355]. While the user may appreciate private defaults, a vendor might prefer to collect data [59]. Recent legislation, such as the GDPR, might place greater responsibility on companies. But when we consider the ‘spirit of the law’, firms would be rational to maximise collection.

Secondly, the frameworks should be applied to systems still in design [355]. Once a platform has been implemented, it is challenging to modify. Smartwatch protection may be constrained by poor usability [183]. However, PbD has little influence now that the interface has been launched. An individual also must have the right and capability to design the system. If Google wished to adjust their designs, they would be free to do so. A privacy-conscious user has little influence over the company’s plans. Finally, PbD addresses system issues rather than human behaviour. We wish

to encourage individuals to protect their data. A platform might have a perfect design, but still be used in an unexpected manner. Therefore, to support privacy awareness, we chose to avoid this approach.

Awareness campaigns

Overview. Many have proposed mitigating the Paradox through increasing awareness [114, 309]. Intuitively, one of the simplest approaches would be an awareness campaign. Such initiatives have been frequently conducted in the UK. Get Safe Online⁵ was established in 2004, and is public-private partnership with the UK government. It provides a website for both consumers and businesses to learn about security. However, its influence has been far from successful. Nine years after the campaign was established, its former chief admitted that it had done “*little to change attitudes*” [236]. This is worrying considering the degree of government support.

Cyber Aware⁶ (previously known as Cyber Streetwise) was launched by the government in 2014. This campaign included a website and adverts in a range of public places. Indeed, with posters displayed in the London Underground, it was likely seen by thousands of commuters. It sought to encourage five basic security techniques, including the use of anti-virus and strong passwords. However, this initiative has also had mixed results. The website cost £12m to develop, but was only accessed by 1.9 million people in its first two years [253]. At a cost of £6.37 per visitor, the approach was not deemed to be cost effective.

Issues. It appears as if awareness campaigns can lack efficacy. This might be true for three reasons. Firstly, highlighting an issue is not sufficient to change behaviour [38]. Campaigns have focused on the threat without providing knowledge or understanding. A person might recognise a problem but have little idea of how to react. Furthermore, if they are purely made anxious, the effort can be counter-productive [359]. Secondly, Sasse et al. [332] suggest that behaviour change should be applied in three stages. Individuals should first be alerted to an issue (‘awareness’). Then education should be provided to improve their knowledge. As a final step, they should be given opportunities to practice. Through this approach, users can repeat this behaviour in future incidents. Finally, awareness campaigns are a large-scale generic technique. This is problematic, as perceptions and attitudes can influence behaviour

⁵<https://www.getsafeonline.org/about-us/>

⁶<https://www.gov.uk/government/news/new-campaign-urges-people-to-be-cyber-streetwise>

[104]. If the information was customised, it might be more persuasive to the recipient. This was a key reason that we selected educational games.

Nudging

Overview. ‘Nudging’ has gained popularity as a behaviour change technique [367]. The approach has been used in governmental units both in the UK⁷ and the US⁸. It is defined as “*any aspect of the choice architecture that alters people’s behavior in a predictable way without forbidding any options or significantly changing their economic incentives*” [367]. For example, banning unhealthy foods might be coercive. In contrast, a nudge could be used by placing salad at eye level. This subscribes to the doctrine of libertarian paternalism, where incentives are used to encourage ‘favourable’ outcomes. Naturally, the definition of ‘favourable’ depends on the person implementing the nudge.

Privacy studies. Due to its popularity, attempts have been made to encourage privacy protection. Wang et al. [384] implemented nudging in a social networking environment. They sought to mitigate unintended disclosure on Facebook. The user interface was modified to both delay posts and highlight their audience. To evaluate the approach, they conducted a 6-week field trial with 28 participants. When the potential audience was made salient, unintended disclosures decreased. Nudges have also been used on smartphones, seeking to encourage the restriction of permissions. Almuhimedi et al. [22] evaluated AppOps, an Android privacy manager, through a 22-day field study. During a treatment period, the tool highlighted the users’ installed applications. After these nudges were introduced, 58% restricted their settings.

Nudging has even been used to address the Privacy Paradox. Jackson and Wang [194] designed personalised notifications for mobile apps. They sought to highlight the discrepancy between a user’s attitude and their current settings. The approach was evaluated by 241 participants in an online simulation. For those receiving the treatment, the disparity appeared to decrease. While this was undertaken in a smartphone environment, the results are highly encouraging. We also adopt dynamically-personalised messages within our educational game.

Issues. However, nudging is still prone to a number of issues. Firstly, the approach can lose efficacy over time, reducing its long-term feasibility [379]. While individuals might be initially influenced, the effect can quickly subside. Although Jackson and Wang [194] appeared to mitigate the Paradox, this was assessed at a

⁷<https://www.behaviouralinsights.co.uk/>

⁸<https://sbst.gov/>

single point in time. If behaviour was evaluated several days later, permissions might have reverted to older configurations. Secondly, nudges can lose their influence once they are removed [70]. In the salad example, behaviour might relapse if the food is rearranged. Therefore, individuals might not be truly ‘persuaded’ to adjust their actions. Again, the Paradox might return when the notifications are removed [194].

Thirdly, implementers must have a right to modify the system. This is often untrue if a vendor has developed the interface. In social networking environments, adjustments can be made through browser extensions [384]. Smartphone nudges are also feasible when the ecosystem is open [22]. However, smartwatches are not amenable to such modification. Although simulations are possible (and used in our Chapter 7 prototype), we wished to analyse empirical behaviour. Finally, nudging has been considered contentious in terms of ethics [217, 318]. The approach is paternalistic, discouraging individuals from certain actions. In contrast, serious games do not weight behaviour or reduce a person’s autonomy. They seek to incentivise an action through engagement and positive reinforcement [101]. For these reasons, we developed educational smartwatch games.

Educational games

These applications were judged to be most appropriate for encouraging protection. They also present great novelty, as privacy games have never been implemented for smartwatch environments. We now introduce the topic, supportive evidence, previous work and the research gap.

Overview. Serious games are defined as “*any form of interactive computer-based game software...that has been developed with the intention to be more than entertainment*” [324]. They have been designed throughout the late 20th century, though their contemporary era began in 2002 [393]. These applications can have several intentions, such as changing behaviour, enhancing understanding or imparting knowledge. ‘Educational games’, also known as Games for Learning (G4L), are a subset that seek to provide teaching [51]. This concept is most relevant to our research, as we aim to increase privacy awareness. Through this act, individuals should be supported in making informed decisions. We hope this will persuade them to protect themselves, thus mitigating the Paradox.

Furthermore, our apps are defined as ‘operative games’, in that they “*leverage knowledge gained from the study of games or play to exert control upon the world such as encouraging exercise or learning*” [79]. We should be explicit that our research does not concern ‘gamification’, which is considered to be a distinct field. This concept

relates to “*the use of game design elements in non-game contexts*” [113]. This might be more akin to nudging, where certain options are made more appealing. In contrast, we develop and evaluate the impact of interactive games.

Success. While serious games have faced scepticism, they have been found to be highly persuasive. Backlund and Hendrix [37] performed a meta-analysis, exploring the efficacy of 40 applications. These games came from a wide range of fields, including language, nutrition and problem solving. They discovered that 72.5% of the implementations yielded positive results. Connolly et al. [101] conducted a systematic literature review on the topic. They examined 7,329 papers and 129 empirical studies. Through this review, they found evidence for the efficacy of games. Wouters et al. [394] also conducted a meta-analysis, comparing serious games with traditional instruction. Their evaluation contained 5,547 participants across 39 studies. Games were found to be significantly more effective for both learning and retention. The field has now gained genuine credibility, with a specialised SIGCHI conference hosted every year [79]. This success gave us confidence that our approach could be persuasive.

Previous work. Security has been frequently addressed by serious games. Sheng et al. [341] evaluated ‘Anti-Phishing Phil’, an application which teaches individuals to evade attacks. The game consists of several levels, in which users have to avoid predators and fraudulent URLs. Phishing detection was evaluated in both pretest and posttest. The game was found to be more educational than tutorials or manuals. Jordan et al. [199] presented CounterMeasures, a game to teach penetration testing. Players were assigned missions and had to complete tasks on a command-line interface. To evaluate the application, participants performed a hacking task at the end of the study. The treatment group were twice as fast as those reading manuals. Raman et al. [316] conducted an evaluation of CyberCIEGE, a 3D simulation. Players were assigned network defence goals, and progressed through the game by completing them. When assessments were performed in posttest, the treatment group received higher results. These findings suggest that serious games might be persuasive.

Privacy research. Despite extensive searching, it appears that only one privacy game has ever been developed. Suknot et al. [363] presented ‘Immaculacy’, a mobile application to highlight the topic. It is a rich interactive story, where characters are immersed in a dystopian world. Points are scored when protective tasks are completed, incentivising the guarding of privacy. We adopt similar techniques in our games, particularly through the use of interactive challenges. Immaculacy also requests responses to troubling scenarios, as found in our questionnaires. However, the influence of this application was never empirically evaluated.

It is curious that privacy games are lacking in popularity. While this provides our applications with great novelty, we have little previous work as a foundation. To address this issue, we adopted Learning Science principles and developed a prototype game. These procedures are both discussed in Chapter 7. There are likely several reasons for the lack of privacy apps. The topic is frequently subsumed within information security. For example, two of CyberCIEGE’s basic topics are ‘information value’ and ‘safeguarding data’ [100]. They could apply to privacy in addition to security. However, the latter field tends to have a greater funding model. While organisations might not pay for privacy, they require their staff to act securely. CyberCIEGE was developed for the US Navy, who would place great importance on this matter. Privacy is also contextual [278], and therefore can be challenging to encapsulate. We seek to address this, as Immaculacy did, through the use of interactive scenarios.

We believe such games could be important for several reasons. Firstly, as previously mentioned, individuals tend to lack privacy awareness [49,247]. It is frequently suggested that the Paradox can be mitigated by increasing understanding [114,309]. Serious games have been found to be highly successful in providing education [394]. Indeed Vaidya et al. [375] recommended gaming scenarios as an approach to teach privacy. Secondly, games allow behaviour to be practised in a safe environment. Privacy settings are daunting [59], and hence users might be reluctant to experiment. If configurations are mismanaged, a person could place themselves at risk. Through providing a simulated environment, individuals can explore without fear. Finally, games might incentivise the usage of privacy protection. The topic is often regarded as a secondary goal [188], and it can place constraints on functionality. However, by supporting and rewarding privacy, users might be persuaded.

Smartwatch research. Few serious games have been created for smartwatch environments. Casano et al. [81] evaluated an app entitled ‘Estimate It!’, which sought to teach measurement and geometry. The game was previously developed for Arduinos, before being ported to a Tizen OS watch. It was assessed by a pool of experts and a 7-person student sample. While users were engaged with the gameplay, the interface had usability issues. Arroyo et al. [31] also developed a tool for the ‘Tangrams Race Game’. It was evaluated through a pretest-posttest design with 96 students. The game was found to contribute to a greater enjoyment of mathematics. However, both these applications were used to augment real-life activities. In contrast, we seek to address behaviour undertaken on the smartwatch.

Others developed smartwatch games, but sought to encourage physical changes. Kim et al. [210] constructed three apps for the Samsung Gear S2. The tools prompted

individuals to engage in stretching exercises. The third application incorporated elements of gameplay to incentivise the action. However, the 42-person evaluation suggested that this was not effective. As demonstrated, serious games are rare on smartwatches. Proposals have been made for ‘Smart Serious Games’, which merge educational techniques with IoT devices [374]. Topologies have also been developed to support this nascent concept [262]. However, smartwatch games remain underexplored, providing our research with great novelty.

We expect this scarcity to be due to several reasons. Firstly, smartwatches are novel devices, having only gained popularity in the last five years. As demonstrated in our Chapter 5 survey, wearables are not possessed by a large proportion of individuals. Therefore, there has been less interest in developing games for this environment. Secondly, smartwatches possess a number of developmental constraints. As suggested by Chun et al. [90], their usability is far from perfect. They have small screens and a limited number of input buttons. Furthermore, due to their computational restrictions, apps could be constrained in their complexity. Finally, educational games are themselves a nascent field. Although the approaches have a rich history, their contemporary era began in 2002 [393]. Other environments, such as desktops and smartphones, currently serve as fertile ground. As serious games gain greater popularity, attention might turn to novel platforms.

We argue that smartwatch privacy games can be of great importance. As previously highlighted, these devices house a variety of sensitive data [118]. Despite this, protective settings appear to be rarely used [371]. Since the environment is novel, we believe users might lack privacy awareness. Without efforts to increase this awareness, individuals might place themselves inadvertently at risk. They might not recognise data collection or protective techniques. Fortunately, serious games have been found successful in providing education [394]. They also have proven efficacy within the field of information security [199, 316, 341]. By instilling knowledge through an engaging platform, we hope to encourage protection. Finally, serious games are most effective when the context is aligned [321]. Since privacy is also contextual [278], smartwatch issues should be addressed through smartwatch apps. Therefore, we develop the first smartwatch privacy game. Through such efforts to provide education and practice, we seek to mitigate the Privacy Paradox.

Chapter summary

In this chapter, we introduced the topics relevant to the thesis. We have also highlighted the research gaps in the Privacy Paradox, Internet-of-Things and smartwatch

games. To describe how these lacunas were addressed, we continue by outlining our consistent methodology.

Chapter 3

Methodology

Overview. In this third chapter, we present an overview of our techniques. This complements the methodology sections included in the subsequent chapters. Since we sought to establish consistency, similar approaches were used throughout the thesis. For brevity, we also describe techniques used on multiple occasions.

This chapter begins by outlining our research philosophy and the resultant approach. Then, to highlight that our measures are appropriate, we justify our selection of quantitative techniques. This is followed by a description of the qualitative measures. We both discuss our inductive analysis [323] and the approaches for validation. Since we sought to evaluate the Paradox in a grounded manner, we continue by presenting our threat models. These were valuable when choosing contextualised questions and protective tools. The selection criteria for these components is also outlined and justified. Finally, as we wished our research to be responsible, we document our ethical procedures.

3.1 Research Philosophy and Approach

Research philosophy

Philosophies. We begin this section by describing our research philosophy. Philosophies can be generally divided into two categories: positivism and interpretivism [175]. Positivists hold an epistemological view which recognises only that which can be scientifically verified [33]. Such individuals seek objectivity, and believe observed phenomena can be rigorously explained. This often leads to quantitative analyses, where metrics can be studied through statistical techniques.

In contrast, interpretivism asserts that individuals must understand the inherent subjectivity in research [382]. Academics should have awareness of how their own

concepts, knowledge and language influence their perceptions. This can conflict with positivism as it posits that multiple interpretations can be equally valid. Interpretivists might prefer qualitative techniques and inductive analyses.

Positivism. For my research, both philosophies appear to have their merits. The Paradox is defined as the disparity between claimed concerns and behaviour. Therefore, to assess the phenomenon, and whether it is mitigated, measurements are necessary. Qualitative methods are informed by the researcher’s interpretation, and hence can introduce subjectivity [346]. Therefore, when comparing concerns and behaviour, we use quantitative metrics.

Our studies also conform to the scientific method by outlining hypotheses and attempting falsification. This differs from interpretivist approaches, which may lack a priori assumptions [71]. Following standards from natural science, ‘control groups’ are used in much of our work. A ‘non-IoT’ sample was included for comparisons in Chapters 5 and 6. Control groups were then components of our final two studies. They are known to mitigate confounding variables, supporting analysis of the treatment groups [344]. Such approaches are deemed to enhance scientific rigour [212].

Our research inhabits the intersection of two fields: Human-Computer Interaction (HCI) and Behavioural Psychology. Since both disciplines take measurements of social phenomena, it has been argued they follow a positivist philosophy [168]. Therefore, our lens appears appropriate.

Interpretivism. However, several elements from interpretivism appear valid. Privacy is a social topic and is therefore open to subjective interpretation [289]. As the concept is nebulous [74], it can be difficult to gauge in an objective manner. Furthermore, since the Paradox can be analysed through several lenses [219], positivism might be challenged.

Due to the contextual influences on privacy [278], we scoped our focus. Such an act introduces the philosophy of reductionism [69]. In this practice, complex phenomena are abstracted into simpler fundamental components. It would be infeasible to evaluate all privacy concepts. Since we wished to explore technological interactions, we analysed ‘information privacy’ [92].

Beyond this reductionism, interpretivism had further influence. Throughout the thesis, we collect large quantities of qualitative data. In the product comparisons (Chapter 5), we solicited rationale for user responses. This was followed by 40 semi-structured interviews in Chapter 6. Although the prototype study was primarily quantitative, we received qualitative feedback from 504 individuals (Chapter 7). Fi-

nally, we conducted detailed interviews in the longitudinal study (Chapter 8). These responses were then analysed through robust techniques.

We opted for inductive thematic analysis, in the style of Ritchie et al. [323]. Rather than applying our own assumptions, findings emerged organically through the rich data. The data-driven approach is described in detail in Section 3.3. Although quantitative measures could produce metrics, they could not dissect rationale. This rationale was crucial to understanding and mitigating the Paradox. Through coding and qualitative validation, we complemented our statistical results.

Chosen philosophy. As previously mentioned, we must assess both concerns and behaviour. This is essential to investigate the Privacy Paradox. To ascertain whether the disparity is mitigated, quantitative measures were used. Therefore, we primarily subscribe to a positivist philosophy. We accept the challenges of quantifying a topic such as privacy [44]. For this reason, concerns and behaviour are contextually-grounded within Chapters 6, 7 and 8. To approach privacy in a nuanced manner, we considered a constructivist lens [338]. However, we feared that the inherent subjectivity might pose practical issues [219]. Nevertheless, since we wish to understand and mitigate the Paradox, we adopt a mixed-methods approach [198].

Methodological approach

Purpose. Through collecting discrete responses, we conduct a large range of quantitative analyses. These techniques, and the justification for their selection, will be outlined in Section 3.2. Such measures allow us to draw consistent results from our collected data. However, while quantitative techniques might tell us the *what*, they do not reveal the *why*. We do not simply wish to confirm the existence of the Paradox. We also seek to understand the disparity (Chapter 6) and mitigate its prevalence (Chapters 7 and 8). Therefore, quantitative comparisons alone cannot address our central question¹. Through a mixed-methods approach [198], we can target the issue.

Approach. In this manner, quantitative results are augmented with qualitative justifications. This methodology offsets the respective weaknesses of the traditional techniques, applying both depth and breadth [198]. It also supports the corroboration of findings, when results reinforce earlier conclusions. For example, the Paradox is confirmed on four occasions: Chapters 4, 6, 7 and 8. Furthermore, we can complement statistical findings with qualitative rationale. In Chapter 7, we assessed participants' concerns and behaviour. Based on discrete responses, we identified that

¹Central question: *Can the Privacy Paradox be mitigated in the context of smartwatches?*

the Paradox was common. However, through qualitative justifications, we discovered the importance of risk perception. By using mixed methods, we sought to ensure that our studies were informed.

Emphasis. Although we use quantitative techniques, our research is directed by qualitative data. Ownership rationale is first explored in Chapter 5. Despite expressed concerns, privacy appeared rarely considered. This encouraged the detailed Paradox evaluations of Chapter 6. Through our semi-structured interviews, we identified the challenges of wearable devices. We also discovered the apparent influence of a lack of awareness. This justified our development of smartwatch privacy games. Qualitative feedback was used to refine the application in Chapter 7. Finally, Chapter 8 explores the rationale behind both concerns and behaviour. Informed by this large quantity of rich data, we seek to both understand and mitigate the Paradox.

3.2 Quantitative Techniques

We now both describe and justify our quantitative measures. Data can be categorised in one of several scales: ratio, interval, ordinal or nominal [358]. We sought to select the most appropriate techniques for each metric. By conducting our analyses in a standard manner, we establish consistency across our chapters. Whenever means are provided, they are denoted by \bar{x} .

Correlation

Spearman. Through our use of Likert Scales, we received a large number of ordinal variables. Therefore, it was inappropriate to analyse correlation through Pearson's Product-Moment technique [294]. Instead, we studied Spearman's Rank-Order Correlation coefficient (r_s) [354]. This was chosen when relationships were assumed to be monotonic, with r_s determining the degree of monotonicity. For example, we used this technique when assessing correlation between privacy concern and disclosure quantity (Chapter 4). We report the coefficient (r_s) and the p -value. p was required to be lower than an α of 0.05 for statistical significance. This threshold is standard for such tests [30]. If true, relationships are probabilistically not due to chance.

Test selection

Scale. We use several measures to assess whether a significant difference exists between groups. Techniques are selected based on three factors. Firstly, we consider

the scale of the dependent (DV) and independent variables (IV). For certain scales, some measures might be appropriate or inappropriate. For example, a Student's *t*-test requires a DV with continuous data (ratio or interval) [361]. If ordinal data is possessed, a Mann-Whitney test would be used instead [250].

Comparisons. Secondly, a measure might be omnibus or two-sample. An omnibus test compares multiple groups to analyse holistic differences. For instance, we might study whether concerns differ based on chosen device (Chapter 5). In contrast, post-hoc tests calculate whether differences exist between individual groups. In the above example, a holistic difference might be discovered. Two-sample tests could then be used to compare two specific devices. When multiple comparisons are made, Bonferroni correction is applied [124]. This process is outlined later in the section.

Groups. Finally, a technique may be chosen for either related groups or independent groups. In the former case, we might have paired results across a passage of time. For example, the pretest-posttest comparisons in Chapter 7. In the latter case, we compare the views of different participants. This can be seen when control groups are studied alongside treatment groups (Chapter 8).

Two-sample tests

Mann-Whitney. When we compared two independent groups, we used the Mann-Whitney U test [250]. This was selected when we had independence of observations, our DVs were ordinal and our IVs were on two nominal levels. For example, it was used when we compared privacy concerns between our control and treatment groups (Chapter 7). We report the *U*-value: the signed frequency at which observations in one group precede the other. The further this value is from zero, the greater the likelihood of a significant difference. We also present effect sizes, which measure the strength of a statistical phenomenon. These are less influenced by sample size than *p*-values [204]. We report Cohen's *d*, with a value of 0.2, 0.5 and 0.8 denoting small, medium, and large effects, respectively [94]. In extreme cases, we highlight 'very large' (≥ 1.2) and 'huge' (≥ 2) effects [333]. We also report the *p*-value and require it to be less than 0.05 for significance.

Student's *t*-test. When our DVs were continuous (ratio or interval), we opted for Student's *t*-test [361]. This was selected if we had no significant outliers and the DVs were normally distributed for each IV group. It was used to compare the disclosure quantity between our public and researcher groups (Chapter 4). As we had independence of observations and homogeneity of variances, an independent *t*-test was performed. We first report the *t*-value: the size of the difference relative to the

variation. The larger this value, the more likely we are to have significance. We also report the degrees of freedom (df): the number of values which are free to vary. The greater this metric, the lower the threshold for significance. As standard, we also show p and Cohen's d .

Wilcoxon Signed-Rank. We used the Wilcoxon Signed-Rank test [391] when comparing related groups. In this case, the same participants responded to multiple queries. This measure was selected if DVs were ordinal, IVs were nominal and the distribution of differences was symmetrical. For example, it was used to compare the privacy concerns towards different products (Chapter 5). We report Z , defined as the number of standard deviations the observed value is above the expected mean. The greater this value, the higher the likelihood of a significant difference.

Chi-Squared. When comparing nominal DVs, we required different techniques. The Chi-Squared (X^2) test [295] was selected when we had independence of observations and a nominal IV. For example, it was used when we compared Paradox prevalence (Y/N) between our two participant groups (Chapter 6). We report four resultant values: X^2 , the degrees of freedom (df), Cramèr's V and the p -value. The first metric measures how well the distribution fits with an expected distribution (if variables were independent). A high value denotes little correlation, implying a significant difference. Since Cohen's d is less suitable for Chi-Squared [105], Cramèr's V illustrated the effect size [105].

McNemar. In the case of repeated measures, we opted for McNemar's test instead [260]. It was used when we had a nominal DV with two mutually-exclusive categories. For example, when we compared Paradox prevalence (Y/N) before and after our prototype session (Chapter 7). We again report X^2 , df , the p -value and V .

Omnibus tests

Kruskal-Wallis. The aforementioned measures are used to compare two sets of data. However, we also conducted omnibus tests, where more than two groups are considered [200]. We selected the Kruskal-Wallis H test [222] when our IVs were nominal, our DVs were ordinal and there was independence of observations. For example, we compared a range of technologies in Chapter 6. To assess whether protective behaviour differed, we made use of this technique. As in the Chi-Squared test, we report the X^2 value, the degrees of freedom (df) and the p -value. Since the technique is omnibus, it was followed by pairwise comparisons. We used Mann-Whitney U tests, before applying Bonferroni correction.

Friedman. When conducting omnibus testing over related samples, we made use of the Friedman test [145]. This is selected when DVs are ordinal and participants are measured more than twice. As explained above, we compared six devices in Chapter 5. To support this, each participant provided six concern evaluations. Therefore, the Friedman test was most appropriate. We report X^2 , df and the p -value. Wilcoxon Signed-Rank tests were then used for pairwise comparisons.

Bonferroni correction. While omnibus tests can signal that a metric varies, they do not highlight the differing groups. To achieve this, we compare each group against its counterparts in a pairwise fashion. Unfortunately, this can lead to the Multiple Comparisons Problem [259]. The more comparisons are made, the more likely significance will owe to chance. We compensate for this by applying Bonferroni correction. This decreases the significance threshold accordingly [124]. For example, we compare the privacy concerns toward six products (Chapter 5). Since this resulted in 15 comparisons, the significance threshold α became 0.003 ($0.05 \div 15$). This approach reduces the risk that findings are claimed when they owe to chance.

Single-sample tests

Occasionally, single-sample techniques are applied on our data. These compare results against a hypothesised distribution, testing whether they differ significantly. If there are three response options, we might expect proportions to be near one-third. However, if one option receives 80%, then the distribution might not be due to chance.

Chi-Squared Goodness-of-Fit. When we have nominal variables, mutual exclusivity and independence of observations, we selected the Chi-Square Goodness-of-Fit test [53]. This was used in Chapter 4 to analyse the proportions of data disclosure. 1% of participants disclosed no data, compared to 72% revealing two elements. Based on this, the proportions did not appear due to chance. We report X^2 , df , Cramèr's V and the p -value.

Single-sample Wilcoxon. When our DVs were ordinal, we chose the single-sample Wilcoxon Signed-Rank test [391]. This was used to ascertain whether values differed from a hypothesised median. For example, we also studied privacy concerns in Chapter 4. We wished to analyse how responses differed from a neutral response. Since 92.8% expressed strong replies, this distribution did not appear due to chance. As for a two-sample test, we report the Z -value, Cohen's d and the p -value. Through the use of appropriate techniques, we consistently analysed our collected data.

3.3 Qualitative Techniques

Data preparation

Purpose. In addition to collecting metrics, we also received qualitative data. Through these extended comments, we can consider the rationale of our participants. Individuals might express several justifications for the Privacy Paradox. By designing responses around these factors, we have a greater chance of mitigating the issue.

However, due to the richness of qualitative data, it can be prone to subjectivity [149]. To extract reliable findings, details must be analysed in a robust manner. Therefore, we conform to the rigorous techniques outlined below.

Formatting. Due to our range of studies, qualitative data was received in a variety of forms. While we collected textual remarks in Chapter 5, Chapter 6 interviews produced hours of audio. Before conducting analyses, the data needed to exist in a standard format. Therefore, as a first stage, we presented our responses consistently.

Most of the responses were textual in nature. These might be received from physical forms or online questionnaires. In the former case, they were transcribed into an electronic document. In the latter case, the responses could be downloaded. They were then removed from the portal for security reasons. The resulting document, usually in CSV format, could then be analysed.

Transcription. In several cases, we conducted semi-structured interviews. We considered collecting responses through rough notes or discrete responses. However, we did not wish to constrain the richness of our data. Furthermore, note-taking might impede the engagement between researcher and interviewee. Therefore, we decided to audio-record these extended discussions. To undertake qualitative analyses, these files required conversion into textual transcripts.

We subscribed to ‘verbatim transcription’ [246], where discussions are detailed in their entirety. Although this approach is time-consuming, it provides the most detailed account of the interaction [246]. Since we wished to dissect behavioural rationale, this was considered important.

Inductive analysis

Overview. For the extraction of qualitative findings, we conducted an inductive process of thematic analysis [323]. This technique, also dubbed ‘framework analysis’, is considered beneficial for four types of research questions: contextual, diagnostic, evaluative and strategic [322]. Since we wish to explore reasons behind a contextual

issue, it appeared appropriate. It is also “*not bound by a particular epistemological position, giving it freedom and flexibility*” and is suitable for varied samples [292].

Data-driven approaches lead to coding which is not placed into a pre-existing model [64]. Although they are often more time-consuming than deductive variants, they seek to avoid bias from latent assumptions [61]. We considered using specialised software (e.g., NVivo) to support our analyses. However, we found coding more flexible when undertaken manually.

Familiarisation. We followed the popular approach outlined by Ritchie et al. [323]. After formatting our data consistently, we began with familiarisation. This was undertaken by parsing through each transcript in detail. The text was read and re-read, enabling initial patterns to emerge. For example, in the prototype evaluation (Chapter 7), we parsed over the responses of our 504 participants.

Annotations. We then began to append brief annotations to the transcripts. These were high-level in nature, and did not seek to be comprehensive. This process was undertaken iteratively as the transcripts were parsed. While a theme might be opaque in the first reading, it might gain salience after further review. Once this process was conducted for all responses, initial patterns began to emerge. Several examples are now discussed below.

In the semi-structured interviews (Chapter 6), participants justified their avoidance of privacy policies. While many found them complex, some were deterred by the effort required. We solicited Paradox justifications in the longitudinal study (Chapter 8). Some cited a lack of awareness, whereas others blamed functionality benefits. These annotations were not final; they merely supported theme construction.

Themes. Based on the annotations, we then inductively developed our themes. This process was undertaken separately for each question. It is possible to encapsulate entire interviews within a set of themes. However, our questions explored a diverse range of topics. To allow responses to be fairly analysed, these queries were considered separately. For example, the views of all participants would be explored together for each question. This approach supports the comparison of concerns, behaviour and rationale between respondents [187]. Themes were formed by reviewing the common annotations. We also included those views which were clearly distinct from other responses. Even if a justification was given by a single person, it was retained. By respecting these deviant cases [23], our themes should represent the data.

The process was undertaken consistently across our studies. In the semi-structured interviews (Chapter 6), there were several themes for password avoidance. While some users considered their data to be innocuous, others thought passwords were

‘too much effort’. In our prototype study (Chapter 7), we asked what would motivate protective behaviour. Some suggested ‘vendor warnings’, while others favoured ‘new helpful features’. Based on these topics, we began to formulate our findings.

Coding indices. When the themes appeared finalised, we produced draft coding indices. These enumerated the themes in a logical order. They also established consistent and representative naming. This skeleton sought to encapsulate the range of responses to each question. A distinct index was developed for each (qualitative) query in the thesis. As previously outlined, this was important due to our diverse questions. When categories appeared too broad, the themes were subdivided to form a hierarchy. For example, there might be multiple reasons a person would oppose data selling (Chapter 6). The individual may either dislike the principle or desire remuneration. To enable a rich analysis, such themes were separated into subthemes.

We also noted the ‘deviant case’ [23], and avoided the temptation to subsume infrequent opinions. Our indices were then adjusted iteratively until we believed the data was faithfully represented. Example indices can be found below in Figure 3.1. The leftmost structure concerns the ownership justifications for wearable devices (Chapter 5). On the right, we present the index for device data access (Chapter 6).

As shown, several themes existed for ownership justifications. These were first divided by each participants’ discrete ‘Yes/No’ response. Within these subthemes, there were many reasons to purchase a wearable. These devices support convenient applications, ranging from notifications to fitness tracking. They also offer advantages in usability and functionality. Similarly, participants had several reasons for rejection. For example, the devices might be considered expensive or prone to breaking. Rationale also varied when considering device data access, as shown on the right. Some individuals feared the reach of hardware vendors, whereas others thought that Google was capable. As before, the hierarchies enabled a detailed analysis of opinions.

Coding frames. Once the indices were finalised, they were developed into our coding frames. These tables can be found in Appendix A. For brevity, we only include frames that are relevant to our textual discussion. Through the use of these tables, we categorised participant responses. To increase the robustness of this process, we developed strict definitions for each theme. While coding, we occasionally noticed that additional themes were necessary. In these circumstances, frames evolved as nodes were added. If ambiguity still existed, the definitions were strengthened to minimise misclassification.

An example frame can be found below in Table 3.1. This encapsulates responses to Question 12 in the Chapter 8 interviews. In this query, we asked participants

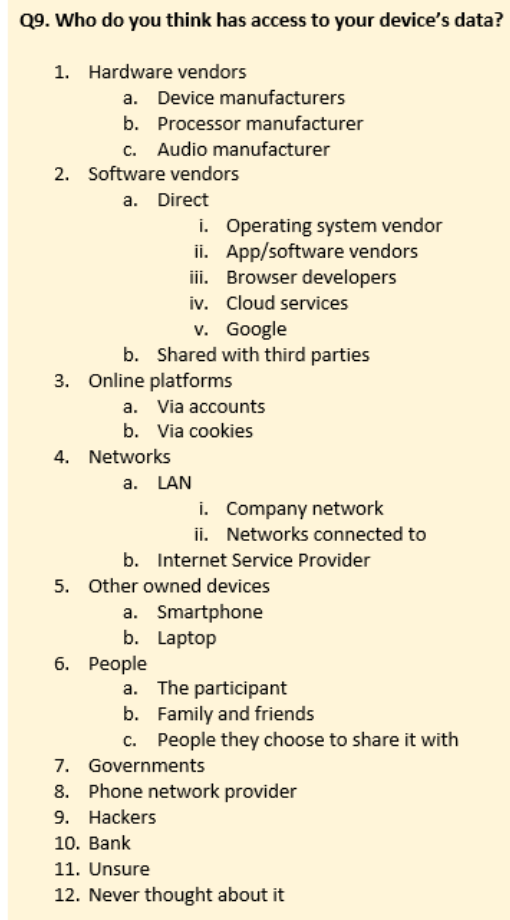
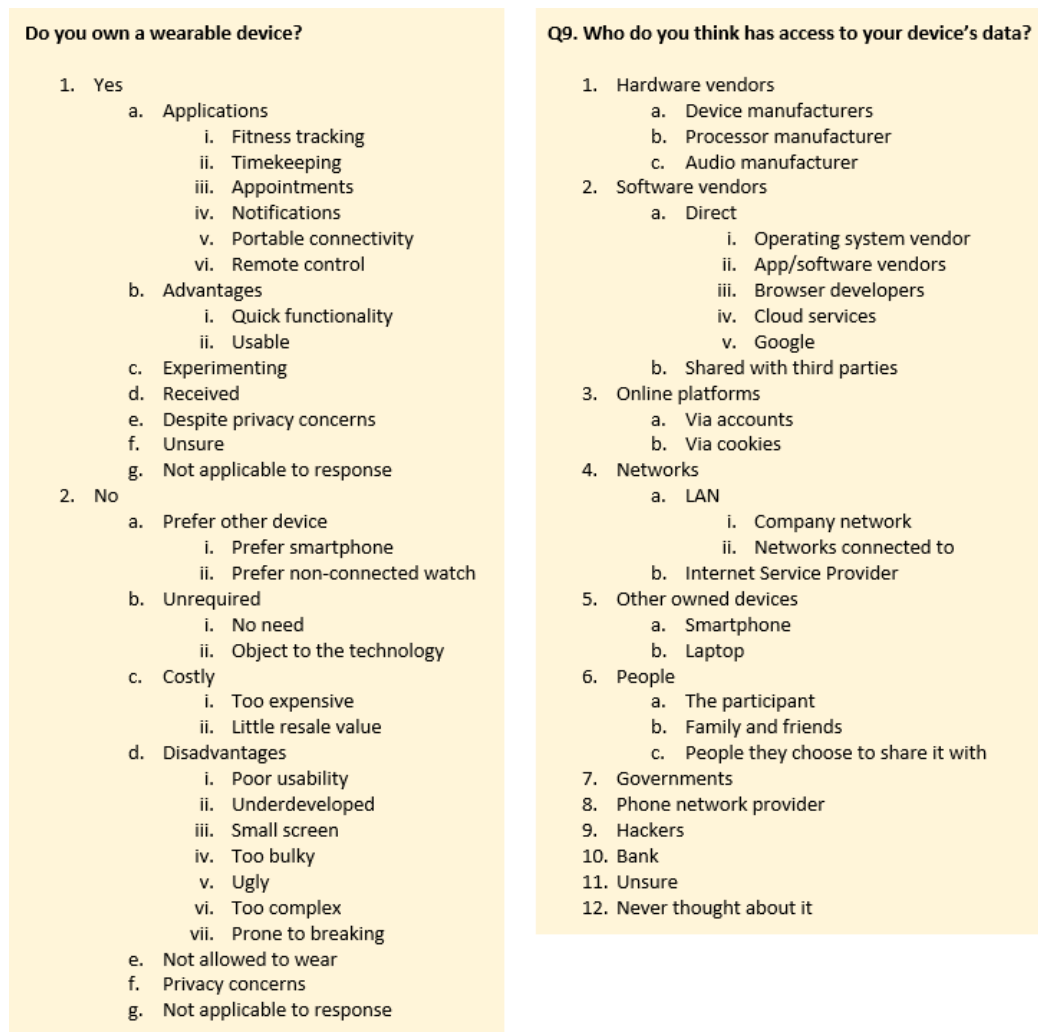


Figure 3.1: Draft Coding Indices: Chapters 5 (left) and 6 (right)

how they might stop apps from reading their smartwatch data. All frames comprise at least three columns. The first represents the top-level themes. These provide a general overview of participant opinions. In some cases, this theme is specified based on discrete responses, such as the reply to a Likert Scale. On these occasions, it is highlighted in the appendices. If subthemes exist, another column is added to divide the category. Similarly, subsubthemes are denoted through further columns. We then include strong definitions for the lowest-level divisions. These seek to reduce ambiguity by distinguishing between topics. The final column houses the relevant responses. For brevity, we only display one quote per theme in this thesis.

Table 3.1: Example Coding Frame: Question 12 in Chapter 8

Theme	Subtheme	Definition	Quote
Approaches	Revoke app permissions	Responses in this category would prevent apps accessing data by changing/disabling the permissions for those apps.	<i>“You could revoke permissions for your contacts or your personal data, etcetera”</i>
	Uninstall apps	Responses in this category would prevent apps accessing data by uninstalling apps which they believe to place their data at risk.	<i>“Or I could uninstall the app if I felt that was needed”</i>
	Avoid/remove sensitive data on watch	Responses in this category would prevent apps accessing data by ensuring they have no sensitive data on their watch.	<i>“Ensuring I don’t have any sensitive data on there”</i>
	Disconnect watch from smartphone	Responses in this category would prevent apps accessing data by disconnecting their watch from the paired smartphone. This would stop the phone from receiving watch data.	<i>“So it’s not connected to my mobile phone”</i>
	Unsure	Responses in this category are not certain on how to prevent apps accessing data.	<i>“Not sure how to go about doing that”</i>
	Impossible	Responses in this category doubt it is possible to delete the data that apps have access to.	<i>“I also don’t think you could make it do that”</i>

Through analysing the frame, we can explore the range of responses. In the aforementioned question, many individuals thought access was limited by revoking permissions. Others believed it was safer to uninstall the applications. Several participants were less sure, or thought protection was impossible. As will be outlined below, a conversation might concern more than one theme. Therefore, the total number of comments might exceed the number of participants.

Coding approach. After the frames were finalised, we began coding our data. This involved copying relevant excerpts into the final column. The same approach was used for all of our qualitative data. Once the transcripts were completed, the frames became fully populated. Through comparing the proportions of themes, we could explore which topics were most prevalent [275].

While themes were independent, a participant might mention several when answering a query. This is commonly found when conducting inductive analyses [307]. To support a rich exploration, responses could be coded with multiple themes. In some cases, topics could conflict or support opposing sides of an argument. For example, we asked individuals about device upgrades in Chapter 6. Interviewees frequently gave reasons both in favour and against this action. As mentioned above, this often meant that responses exceeded the number of participants. For this reason, we made a distinction between two proportions: *comments* and *participants*.

A *comment* was defined as an utterance that has been coded. An individual is likely to mention several codes in a single answer. Each distinct remark would be considered a separate comment. Therefore, if a person repeatedly made the same point, it was only counted once. The total number of comments is the sum of all the codes mentioned by all participants. In the above frame, 16 distinct remarks were made. However, only 10 individuals took part in the study. Since ‘uninstall apps’ was mentioned on four occasions, this would result in a percentage of 25% (4/16).

In the case of *participants*, we study how many individuals cited a certain code. For example, four out of 10 users spoke of uninstalling apps. This would result in a participant percentage of 40% (4/10).

Reporting. Once coding is complete, we receive a collection of populated frames. We can then extract findings through several approaches. Firstly, we highlight those perceptions which are most prevalent. For example, in Chapter 6 we explored the factors which influence the Paradox. Since a ‘lack of awareness’ was the most common theme, it appeared a sensible candidate. Secondly, we investigate interesting but uncommon responses. While these remarks challenge categorisation, they provide unique insights. For example, in Chapter 5, one person bought a tablet despite privacy concerns. This might suggest that functionality was considered more important.

Thirdly, we compare perceptions between participant groups. If their theme proportions are dissimilar, it suggests that viewpoints are different. Such an approach is frequently used to extract findings from qualitative data [275]. For example, we compared our treatment group and our control group in Chapter 8. When discussing unauthorised access, treatment individuals feared for their personal data. In contrast, 40% of others doubted their watch’s sensitivity. Finally, when presenting results, we highlight vivid examples [64]. These are participant remarks that typify a specific theme. Throughout this thesis, they are displayed alongside qualitative results. By grounding our findings in rich data, our conclusions should be true to our participants.

Qualitative validation

Overview. Through our inductive approach, we undertook a detailed analysis of our collected data. However, due to the richness of our content, there were risks of misinterpretation. To add further rigour to our process, we then performed qualitative validation. This was achieved through four techniques: *triangulation* [139], *repeated coding* [245], *multiple coding* [293] and *respondent validation* [56]. Due to practical constraints, each approach was not used in every study. However, they introduced rigour when it was most required. We outline our techniques below.

Triangulation. In this process, we receive data from multiple sources. Then, by comparing these results, we can corroborate our findings. This mitigates bias and variation found in the individual sources [139]. If we receive a result in one experiment, we can have a degree of confidence. However, if this is supported by a different instrument, the conclusion is more likely [139]. At a high level, our thesis uses frequent triangulation. For example, the Paradox’s presence is confirmed in Chapters 4, 6, 7 and 8. This gave us confidence that the issue is prevalent. Privacy rationale is also explored through multiple studies (6, 7 and 8). Since justifications appeared to align, we are confident in our findings.

Triangulation: Process. Within individual chapters, triangulation is also used. In Chapter 4 we do not collect qualitative data. However, behaviour is studied through both disclosure and online actions. In both cases, behaviour rarely appeared to be privacy-protective. This seemed to encapsulate the position of our sample.

Our semi-structured interviews (Chapter 6) explored a variety of devices. To increase the richness of our data, topics were approached from several angles. For example, we asked participants why privacy-invasive products are used. We also requested justifications for the Privacy Paradox. In both cases, a lack of awareness was frequently mentioned. This gave us confidence to target the issue.

We conducted our longitudinal study in Chapter 8. To increase rigour, we triangulated through using both questionnaires and interviews. In the surveys, we explored participants’ responses to privacy incidents. We then discussed potential risks in the interviews. Since rationale aligned on both occasions, we were confident in our results.

Repeated coding. We sought for our coding to be as robust as possible. This was undertaken by refining frames and using strong definitions. However, when data is coded once, subjectivity might have some influence [346]. This can occur for several reasons. Responses might not directly align with themes, challenging the categorisation. Alternatively, the phrasing of the data might lead to ambiguity. To increase rigour, coding can be conducted twice over several weeks. By comparing the results, one can calculate their *intra*-rater reliability [245]. This can be used to identify frames in need of refinement.

Repeated coding: Process. The longitudinal study (Chapter 8) was the apex of our work. Since it addressed our central question², we sought to validate its results. Therefore, the doctoral researcher performed repeated coding on all the data. Although this was time-consuming, it presented a chance to refine our frames.

²Central question: *Can the Privacy Paradox be mitigated in the context of smartwatches?*

Once the first coding was complete, we populated a framework matrix [147]. For illustration, the gameplay matrix from Chapter 8 is found below in Table 3.2. Columns represent participants, while rows denote questions. Each cell contains the coded theme, with more than one often assigned. For example, Participant *J* appreciated both the interaction and performance of their game.

Table 3.2: Example Framework Matrix: Chapter 8 Gameplay Responses

Questions	Treatment Group					Control Group				
	A	B	C	D	E	F	G	H	I	J
Like	Layout	Security/privacy Interaction	Character	Simplicity Aesthetics Character	Security/privacy	Interaction Aesthetics	Flexibility	General Interaction Simplicity	General	Interaction Performance
Not like	General Navigation Too-difficult	Repetitive	Challenge- variety	General Navigation Screen-lock- challenges	Screen-lock- challenges	General Navigation	Theme Repetitive Too-easy	Repetitive	Map-variety	Repetitive
Improve	Challenges Questions Enhance-interface Reduce-challenge- difficulty	Levels Difficulties	Fix-challenge- bugs	Fix-challenge- bugs	Challenges Questions Levels	Challenges Enhance- interface	Questions Battle-system	Challenges	Greater-maze Level-menu	Challenges

After three weeks had passed, we performed the second coding. This period was installed to reduce retention from the first parse. The same frames were used, since we sought to validate and refine their structure. Another matrix was then populated with the new results. We analysed consistency by comparing the theme distributions through ‘proportion agreement’ [271]. In the above case, 54 themes were assigned across both rounds. If 2 were unique to the first and 2 were unique to the second, we would match on 50/54 occasions. This would result in an agreement rate of 92.6%. We selected this method over Cohen’s kappa [93] for two reasons. Firstly, there were a large number of themes, reducing the risk that matching is due to chance. Secondly, since responses often mentioned multiple themes, kappa is not appropriate [93].

Since the results have relevance to Chapter 8, they are outlined in Section 8.4. We considered using repeated coding for our Chapter 6 interviews. However, with 40 discussions of 20 minutes, the time investment was deemed excessive. We thought this time was better spent on validating the final study. In both Chapter 5 and Chapter 7, we received quantities of qualitative data. However, since responses were brief and textual, a single parse was felt sufficient.

Multiple coding. Researchers can often judge a matter in different ways. Provided with the same data and coding frames, individuals might produce several distributions. Only if the frames are precise, will the results be consistent. To introduce further rigour into our analysis, we undertook multiple coding [293]. In this process, two individuals code the same transcripts. They are provided with frames which they gradually populate. Once complete, their theme proportions are compared. If the

results are similar, it gives confidence to the analysis. However, if the proportions differ greatly, the coding is prone to excessive subjectivity.

This approach is recommended as good practice in qualitative research [293]. Since our longitudinal study was of prime importance, its data was multiple-coded. While we considered the technique for other chapters, the process was not deemed necessary. In the product perception surveys (Chapter 5), qualitative data was short and brief. Similarly, responses were not extensive in the prototype study (Chapter 7). This was accepted, as our online approaches balanced richness against sample size. Our semi-structured interviews (Chapter 6) were a candidate for multiple coding. However, with 40 discussions of 20 minutes, it would have taken considerable time. Furthermore, we lacked another student for consistent collaboration. Therefore, the time was deemed best invested into the longitudinal study.

If two individuals undertook inductive analysis, they would each approach the topic with different viewpoints. Therefore, it is improbable that their coding frames would be similar. In our studies, we did not wish to receive a different interpretation. Since our frames were informed through five studies, the lens was appropriate. Instead, we sought to validate the robustness of our coding. To achieve this, the second researcher populated our frames. The agreement rate was then used to highlight areas for refinement.

Multiple coding: Process. We now briefly outline our process for this multiple coding. The first parse was undertaken by the doctoral student. They used the refined frames from the repeated coding (as outlined above). Based on the proportions of each theme, we then populated a framework matrix [147]. As before, columns represented participants and rows denoted questions. We then approached a second coder, who was another member of our research team. They had recently joined the university and were not familiar with the student's work. They also lacked a Computer Science background, unlike the student. Their coding offered a rigorous evaluation of our frame's comprehensibility.

To begin, the second coder was given the interview transcripts. They also received the coding frames, with one created for each question. The excerpts were removed from the tables, ensuring their categorisation was not influenced. From the questions, theme names and definitions, they repopulated the frames. Once complete, we produced a framework matrix from their work. To compare the theme distributions, we again opted for proportion agreement [271]. This approach was selected for the aforementioned reasons. Since the results are more relevant to Chapter 8, they can be found in Section 8.4.

Once the comparison was complete, we considered the areas of disagreement. Their definitions were updated to reduce further ambiguities. This produced the final frames, which were then used for our qualitative analyses. Since we followed this rigorous process, we believe our research was undertaken in a robust fashion.

Respondent validation. Finally, we outline the technique of respondent validation [56]. Researchers might be in agreement over an individual’s responses. However, the participant is not involved in this process. Therefore, there is a risk that their views are misinterpreted. Findings could be based on transcription errors or ambiguities. This is particularly likely when transcribing large quantities of audio. However, if we liaise with participants, we can verify whether we have captured their opinions.

Respondent validation seeks to involve these individuals in the process [56]. To achieve this, we assigned our interviewees a number of tasks. Firstly, they verified the accuracy of their transcript. This reduces the risk that simple errors are introduced. If the textual account is accurate, we have a solid foundation for our analysis. Secondly, since this approach follows coding, they checked the appropriateness of their assigned themes. If an answer was ambiguous or misinterpreted, the topics might not match a participants’ views. However, as is standard in respondent validation, amendments were not automatically honoured [56]. Although interviewees know what they said, they often lack subject-matter expertise. For example, while a participant might regard themselves as informed, they may express a low degree of privacy awareness. In this case, they might not recognise the threat or protection. Therefore, requests for amendment were weighed against our own judgement.

These two tasks were completed for our Chapter 6 interviews. For our longitudinal study (Chapter 8), an additional stage was included. Based on responses to Protection Motivation Theory questions, we produced several ‘participant profiles’. These comprised textual characterisations of each person’s rationale. To inform our findings, it was important that these profiles were representative. Therefore, interviewees were also asked to check the descriptions.

Respondent validation: Process. A similar approach was undertaken for both studies. After the discussions were completed, we sent each participant their interview transcript. This transcript had been reviewed several times to minimise the error rate. Beneath each answer, the relevant codes were listed clearly. An example transcript can be found below in Figure 3.2. This excerpt illustrates an answer to Question 8 in Chapter 6. The main query is highlighted in blue, while follow-up questions are left-aligned. Participant responses are indented, with the assigned themes shown in bold. When multiple subthemes are referenced, they are listed within the brackets.

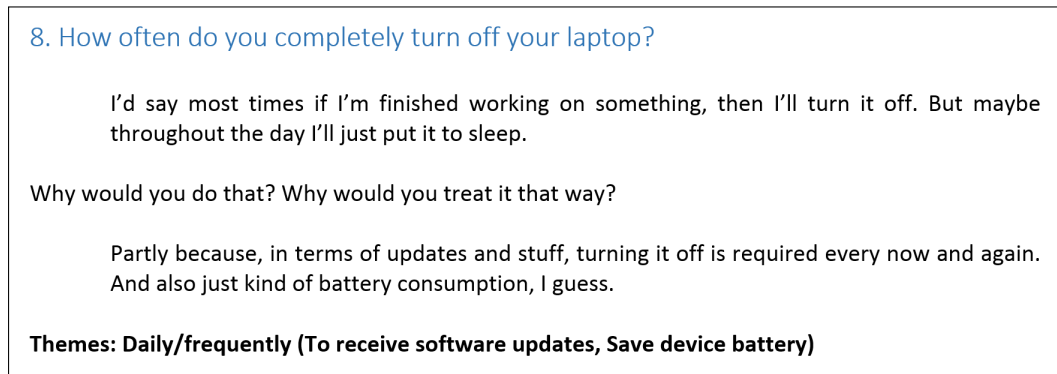


Figure 3.2: Respondent Validation: Question 8 in Chapter 6

In both studies, interviewees were given two initial tasks. They first verified the textual accuracy of the transcript. They then assessed the appropriateness of the assigned codes. This judgement had to be made based on the recorded responses. This prevented a participant from changing their rationale after the interviews. In the longitudinal study (Chapter 8), interviewees also checked the accuracy of their descriptive profile. Once these tasks were complete, participants responded to the researchers by email. They could request amendments on any of the aforementioned grounds. The specific results have greater relevance to the chapters. Because of this, they are detailed in Sections 6.3 and 8.4, respectively.

Summary. As outlined above, we use four techniques for qualitative validation. Triangulation seeks to corroborate findings through data from multiple sources [139]. This approach was used in Chapters 4, 6 and 8. At a high level, triangulation is also used throughout the thesis. We undertook repeated coding to both validate and refine our frames [245]. The approach is found in Chapter 8, since this study addressed our central research question.

Multiple coding is a similar technique, but requires the participation of an independent researcher [293]. As it seeks to introduce greater rigour, it was also used in our longitudinal study. Finally, once coding was complete, we implemented respondent validation [56]. Through this technique, respondents verify that they have been appropriately interpreted. Since interviews are most prone to ambiguity, it was used in Chapters 6 and 8. By adopting validation, we seek to conduct robust analyses.

3.4 Threat Models

Overview

Purpose. In Section 3.5, we describe our process for selecting contextualised questions. These are used throughout the thesis to gauge privacy concerns. The queries solicit a Likert-scale response to a troubling incident. This follows the popular approach of Lee et al. [232]. These authors explored concerns on wearables, and hence the technique appeared fitting for IoT devices. While they labelled their queries as ‘scenarios’, we opt for ‘contextualised questions’. This avoids conflating the approach with that of descriptive vignettes [135]. However, concern queries are still grounded within a rich environment. This supports a nuanced analysis of participant opinion.

As will be highlighted in Section 3.5, the questions had three selection criteria. Of central importance was *feasibility*, since it is crucial that the incidents are possible. If issues are infeasible, we learn little from participant unease. To make such a judgement, we first required an explicit threat model. This construct outlines the risks that are deemed in scope within a particular context [274]. To support grounded analyses of concern, we selected sensible models.

Overview. Our research requires two distinct threat models. At the start of this thesis, we explore privacy in a general sense. Although we scope to ‘information privacy’ [92], a range of devices are considered. This allows concerns to be compared in the Chapter 6 interviews. To gauge perceptions across several products, we need a *general* threat model.

However, we refine our scope as our research progresses. The interviews suggest that the Paradox is common in wearable environments. Therefore, smartwatches are explored in Chapters 7 and 8. The first threat model would be inappropriate for these analyses. There are risks, such as device theft, which might be more common to wearables [320]. Furthermore, certain threats might have lesser relevance to smartwatches, such as malware [213]. As a result, the *smartwatch* model was also developed.

General model

Considerations. For the general model, we considered high-level risks to common devices. Since we analysed a range of products, it was required to be broad. If not, incidents facing some technologies might be considered infeasible. This would bias the comparisons of privacy concern. We believe many individuals would desire greater data protection. However, this eventuality is impeded by several factors. Firstly, since online privacy tends to lack salience [3], it is often forgotten. Secondly, as

settings are frequently opaque or hidden, protection might be challenging to achieve [59]. Finally, privacy is usually a secondary goal [188]. Because of this, users might become preoccupied with device functionality. Based on these considerations, we constructed our first threat model.

Threat model. Individuals were assumed to have some interaction with technological devices. These products could range from desktop computers to wearable gadgets. Through this interaction, they would receive functionality benefits. However, they might also encounter a number of risks. Cybercrime is prevalent, with £130bn lost globally in 2017 [173]. Attacks can take many forms, from hacking to phishing to malware infection. Malicious parties might infiltrate a computer to access its data. They could also delete information if the machine becomes compromised. These issues are increasingly common, as seen in last year’s ransomware epidemic [172]. Data might also be desirable to technology companies. Depending on settings, details can be accessed through apps or operating systems. Although access is often optional, users rarely read their privacy policies [155].

Free services are provided by several large corporations (e.g., Google, Facebook, etc.). Due to their popularity, we assume they are used by most individuals. These services might possess privacy features, such as passwords, permissions or policies. However, such settings often lack salience or usability [197]. To fund free services, companies frequently collect consumer data [335]. This has been demonstrated by the privacy controversies in 2018 [156]. Collected details can then be used in a variety of ways. Data be deleted or processed no further. It could also support marketing and advertising efforts. Alternatively, it could be shared or sold on data exchanges [334]. These practices are common and support much of the digital economy [334]. Therefore, if a user encounters a privacy risk, it likely originates from one of two sources: tech companies or malicious parties. The former case may be consented, while latter issues are unconsented. Since these incidents are reasonably common, we believe this to be an apt threat model.

Smartwatch model

Considerations. After the Chapter 6 interviews, we scoped our focus to smartwatches. As justified in Section 7.2, we selected the (Android) Wear OS ecosystem. Therefore, we required a distinct threat model. Unlike in the previous case, individuals are unlikely to use web browsers. Since the market is smaller, there might also be less targeted malware [213]. However, users do operate within a constrained interface. Functionality might also be dependent on third-party applications. Furthermore, as

GPS is possessed on many watches, positions might be identifiable [32]. Due to these factors, users are likely to face different threats.

Threat model. Individuals are assumed to use a smartwatch. Their device can support convenient applications, which are installed through an app store. Many of these programs are developed by third-party vendors. These companies might have an incentive to access smartwatch data. While data supports personalisation, many firms are funded by processing these details [59]. This data could then be shared or sold (potentially aggregated) to external parties [334]. Fortunately, smartwatches possess settings to restrict collection. By configuring app permissions, the access of companies can be limited. However, since this approach can reduce functionality, it is often overlooked by users [36].

Location services are supported, whether natively or through a paired smartphone. GPS offers useful features, such as exercise tracking and route navigation. However, to provide this functionality, the device’s current location is identified [32]. Although locations can be inferred through WiFi or mobile networks, GPS provides a precise position [398]. Smartwatches are also small, expensive, popular and consumer-oriented. Like smartphones, this places them at a reasonable risk of loss or theft [256]. Indeed, as expressed by Ricci et al. [320], their “*size and portability makes them easy to steal*”. Since watches are novel and visible, they might attract attention. Therefore, based on the above points, threats might arise from app companies or petty criminals.

3.5 Consistent Selection Criteria

Concern contextualised questions

Purpose. To analyse the Paradox, we must explore participants’ privacy concerns. However, as previously mentioned, the principle is inherently contextual [278]. Furthermore, when privacy is primed, individuals’ concerns are often inflated [65]. Therefore, the topic is challenging to assess in a fair manner.

Scenarios have proved effective in many previous studies [2,133,232]. The approach grounds privacy risks within a particular context. By evaluating reactions to these incidents, we can ascertain the degree of concern. This avoids the term ‘privacy’, also reducing the likelihood of priming. Lee et al. [232] adopted this approach when analysing wearable concerns. Since we explore novel environments, we believed it would be appropriate. However, unlike Lee et al., we do not label our queries as privacy scenarios. We opt for ‘contextualised questions’, as this avoids conflating the

technique with descriptive vignettes [135]. To support rich evaluations, we use the approach in Chapters 6, 7 and 8.

Introduction. Contextualised questions must be carefully constructed. If an incident is irrelevant then it serves little purpose. Similarly, if the issue is infeasible, concerns are not truly assessed. To select scenarios, we followed a consistent methodology. In Chapter 6, the process was used to choose high-level incidents. This was important since we compared perceptions across a range of devices. In Chapter 7, the issues were specific to smartwatches. In Chapter 8, the incidents directly concerned our selected model (Huawei Watch 2). As the same selection criteria were used in all cases, we established consistency throughout our research.

The criteria comprised: *relevance*, *feasibility* and *behaviour correspondence*. The rationale for each criterion is provided below.

Relevance. Incidents should all be relevant to the topic of privacy. Otherwise, expressed concerns would relate to other topics. To increase the chance of relevance, questions were matched to Solove’s archetypal violations [352]. In his popular taxonomy, he outlined the methods in which privacy is invaded. These include intrusion, surveillance and secondary use. If scenarios relate to such violations, they are likely to concern privacy. Our research scope should also be noted. We consider ‘information privacy’, rather than the physical varieties [92]. Therefore, issues should have relevance to digital devices.

Feasibility. For a fair assessment of concern, the incidents should be feasible. If a violation is devastating, then participants will clearly express unease. However, if that threat is impossible, concern evaluations tell us little. To gauge feasibility, we considered our two threat models. Since we compared several devices in Chapter 6, we chose the *general* model. In Chapters 7 and 8, we used the *smartwatch* construct. If the issue was possible within our sensible models, it was deemed to have feasibility.

Behaviour correspondence. To explore the Paradox, behaviour is compared with concerns. For a well-grounded comparison, actions should mitigate the scenario risk. If a setting can reduce a feared threat, one would hope it is used. Therefore, incidents must be addressable through available protection. In Chapter 6, these settings should be common to a wide range of devices. For Chapters 7 and 8, they should be common to smartwatches. This establishes the *principle of compatibility*, where questions relate to the same topic [12]. Such a correspondence is deemed to support appropriate comparisons [12].

Selection. The chosen questions are outlined in the individual chapters. Their selection is justified based on adherence to the above criteria. Through this consistent and robust approach, we seek to evaluate concern in a grounded manner.

Privacy-protective features

Purpose. We now consider the protective features. These are defined as settings which can be used to enhance privacy. Their selection also followed a robust and consistent process. This was important for our analyses of participant behaviour.

Introduction. We selected features through four criteria. These comprised: *simplicity*, *utility*, *applicability* and *concern correspondence*. Features were considered in Chapters 6, 7 and 8. In the first case, the settings were required to be general. Since we evaluated a range of devices, the features must be common. Our focus was then scoped in the latter two chapters. As they analysed smartwatches, the features were selected from Wear OS settings.

Simplicity. It should be feasible for all tools to be used. A person might possess genuine concerns about their privacy. However, if protection is excessively complex, they have few opportunities. To support a fair assessment of behaviour, we took two approaches. Firstly, tools should be described in an accessible manner. To achieve this, we avoided the usage of technical terminology. Secondly, settings should not require Computer Science expertise. Therefore, we excluded tools with command-line interfaces or USB configuration.

Utility. For a feature to have relevance, it must provide some benefit to privacy. Devices can possess a wide range of settings. However, if a tool provides no assistance, its usage is irrelevant. Therefore, we required features to either raise privacy awareness, increase control or protect data. When settings offer such benefits, user privacy is supported.

Applicability. Chosen features must be widely available on the selected devices. If the settings are rare, then their usage will be uncommon. However, this would be due to unavailability rather than lax behaviour. As previously mentioned, several devices are considered in Chapter 6. The features must be supported across the range of studied technologies. Smartwatches are analysed in Chapters 7 and 8. In the former case, we consider these devices as a whole. In the latter case, we specifically explore the Huawei Watch 2. For well-grounded evaluations, settings must be common in their respective environments.

Concern correspondence. As stated in the contextualised question criteria, settings should mitigate the risk from incidents. If not, participants cannot be criticised for inaction. Often the issues were chosen after the protective features. In these cases, there would be no incidents for us to consider. To address this, settings should be selected which mitigate a range of risks. Then, when scenarios are chosen, they can establish ‘behaviour correspondence’. By connecting concerns and actions, we respect the *principle of compatibility* [12]. This supports rich assessments of the Privacy Paradox.

Selection. The selected features are described in Chapters 6, 7 and 8. In each case, the decision is justified in reference to these four criteria.

3.6 Ethics

Purpose. Since we conduct research on human behaviour, ethical considerations were essential. We have a duty to undertake studies in a responsible manner. This is particularly true since we interact with members of the public. Furthermore, participants might be discouraged if ethics are not considered. This could lead to non-response bias [164], with privacy-conscious individuals avoiding our research. To minimise these risks, we followed the below procedures.

Ethical approval. Before each study was conducted, we submitted our designs for ethical approval. All approvals went through our institution’s Central University Research Ethics Committee (CUREC). As part of this procedure, a wide range of documents were reviewed. These included methodologies, questionnaires, interview queries and consent forms. After evaluating this information, the committee could suggest amendments or seek clarifications. Once agreement was met, the research received ethical approval. This was achieved for all the work in this thesis.

Informed consent. Consent was requested before individuals participated in our studies. This was achieved through several processes. For our street surveys (Chapter 4), an oral description was most appropriate. This was due to the brevity and informality of the interaction. However, procedures were also included on the questionnaire forms. Participants signed this document if they agreed with the conditions. Since we distributed questionnaires at several sites, we also received authorisation from the local councils. In Chapter 5, we compared product perceptions through an online survey. As the respondents were remote, the terms were included in the digital portal. An online approach was also used in Chapter 7, where smartwatch owners evaluated

our prototype game. Since they participated in three stages (pretest, gameplay and posttest), they were asked for consent on three occasions.

In Chapters 6 and 8, we directly interacted with our participants. In these cases, they were provided with written information sheets. The sheets fully outlined research conditions and data protection. Again, individuals signed the forms to indicate agreement. By acquiring consent in each study, we sought to conduct our work responsibly.

Chapter summary

In this chapter, we outlined our high-level methodology. We also justified the usage of a range of quantitative and qualitative techniques. Since concerns are gauged through contextualised questions, we described their selection. After discussing protective settings, we finished by highlighting ethics. Through subscribing to a standard process, we establish consistency throughout the thesis. With the methodology outlined, we now proceed to our first study.

Chapter 4

Can the Privacy Paradox be confirmed in the UK?

4.1 Introduction

Background

Previous work has highlighted the disparity between privacy concerns and behaviour [272]. These Paradox studies have been conducted in a range of environments, from web browsing [207] to smartphone use [298]. However, while many works are conducted in the US [4, 122, 148, 163, 279, 356], the issue is underexplored in other countries. As previously mentioned, there has been some research in the UK. Brown [68] was among the first to uncover the Paradox, studying perceptions of online shopping. Unfortunately, although he interviewed UK residents, the sample size was small ($n = 12$). At this point the Paradox was suggested, rather than carefully confirmed. Zafeiropoulou et al. [397] assessed the matter through a 150-person survey. They found that privacy decisions are influenced by external factors such as location. However, despite these insights, less than half of their participants came from the UK. The Paradox remains underexplored in this context. Privacy is inherently cultural [18], with perceptions of the topic differing across the world. This could result in a variety of concerns and protective behaviour. Therefore, before considering the IoT, it was wise to confirm the Paradox's existence in the United Kingdom. If not, we might conduct studies on an absent phenomenon. Furthermore, without a solid foundation, IoT analyses lose their relevance. Once a concern-behaviour disparity is confirmed, we can then begin deconstructing the issue.

Motivation

As will be detailed, we targeted members of the public through a street survey. This allowed opinions to be solicited in an accessible environment. The survey was undertaken in four distinct locations: London, Birmingham, Oxford and Cardiff. Since we wished to analyse a range of views, a single-site study would have been inadequate. These cities all differ in terms of region, population size and ethnic demographics. As perceptions might vary based on background, this offered a chance to survey a spectrum of opinions. We do not claim to be truly representative of the public, as few studies can. Indeed, as outlined in our critique (Section 9.3), no sites were based in Scotland. However, if the Paradox is prevalent in our sample, some UK residents should be susceptible.

While this study sought to confirm the Paradox, we also had an opportunity to explore contributory factors. In previous work, a lack of awareness has been commonly highlighted [114,309]. If a person does not understand privacy, they might not recognise the risk [48]. Furthermore, if they lack knowledge of protection, they will have little chance to guard their data. To explore the influence of privacy awareness, we compared public responses to those of an expert group. This was undertaken through recruiting a distinct sample of cyber security researchers. As these individuals had knowledge of privacy tools, we expected their behaviour to be protective. If so, it might suggest that awareness can encourage protection.

Methodology selection. A street survey was preferred over alternative methodologies. We considered beginning our research with a comparative study. This would compare perceptions of IoT and less-novel technologies, much as is detailed in Chapter 5. However, without first confirming the disparity, we believed it an unwise approach. We also considered using an online questionnaire for this UK study. Such surveys are cheaper to conduct and have a lower time investment [192]. Nevertheless, we decided that our street approach was preferable. Through conducting the research in public, we could access demographics which are underrepresented on the Internet [24]. By distributing the forms in different sites, we also increased the chances of varied opinions. Furthermore, as online questionnaires were used in later studies, this offline approach could complement their findings.

4.2 Methodology

Hypotheses

We begin the methodology section by discussing our hypotheses. These assertions directly contribute to our central research question. Reiterating from Chapter 1, this was: *Can the Privacy Paradox be mitigated in the context of smartwatches?* To make progress in answering this question, we must first address a more pertinent subquestion: *Can the Privacy Paradox be confirmed in the UK?* For an exploration of this matter, we extracted both respondents' concerns and behaviour. We also wished to explore the influence of privacy awareness. To achieve this, we must compare public views against those of our researcher sample.

Metrics. Through our survey, we evaluated four distinct factors. Two related to perceptions: *privacy concerns* and *self-perceptions*, while two regarded actions: *online behaviour* and *empirical behaviour*. We begin by discussing privacy concerns. To assess whether individuals were concerned about the topic, we solicited whether they considered online privacy to be important. This sought to match our scoped focus of 'information privacy' [92]. While we considered referencing 'concern' explicitly, we thought this might inflate responses [65]. Participants answered on a five-point Likert Scale ranging from 'Strongly Agree' to 'Strongly Disagree' via 'Neutral'. This ordinal response then produced the *concern score*, which spanned from 1 (low) to 5 (high). If a person is concerned about privacy, they should find it important.

We also explored how individuals perceive their own actions. These *self-perceptions* can reveal the behaviour that participants deem appropriate. To assess perceptions, we asked respondents whether they were private with their own data. Again, we used a five-point Likert Scale ranging from 'Strongly Agree' to 'Strongly Disagree' via 'Neutral'. This response comprised the *perception score*, spanning from 1 (low) to 5 (high). If participants consider themselves to be privacy-protective, this suggests concern for the principle. In contrast, if they believe they act poorly, it might signify a lack of interest.

After examining concerns, we moved forward to consider privacy-protective actions. They were explored through two factors: *online behaviour* and *empirical behaviour*. To assess the former, we solicited the usage frequency of several protective measures. These included app permissions and data encryption, with the full list outlined in the Survey Design. We selected a four-point Likert scale, ranging from 'Always' to 'Never' via 'Often' and 'Rarely'. This was chosen since it was unlikely

that measures would be used consistently. To compose the *behaviour score*, we computed a mean frequency across our 10 techniques. It ranged from 1 (low) to 4 (high). This was deemed appropriate, since the responses were aggregated around the same topic [58, 78, 280]. If an individual takes frequent action, then they express privacy-protective behaviour. Since the public might lack awareness of protection [49], we expect their usage to be rare. In contrast, we believe that researchers will possess the expertise to use these tools.

In addition to online behaviour, we sought to analyse empirical actions. This analysis was undertaken subtly to minimise the risk of response fabrication. Our questionnaire requested seven demographic details, three of which were optional. These fields included a ‘Prefer Not to Say’ option and were explicitly highlighted as optional. To explore *empirical behaviour*, we assessed the degree of disclosure by each individual. While not a perfect proxy for protective actions, it is indicative of their likelihood. Based on the number of revealed details, we composed the *disclosure quantity*. As before, the questions will be outlined in the Survey Design. Since it is popular to share data, we expected our public sample to frequently disclose. Researchers should be aware of privacy risks, and therefore might exercise greater caution.

This design resulted in four metrics: *concern scores*, *perception scores*, *behaviour scores* and *disclosure quantity*. Through studying these factors, we explored six hypotheses. The first two assessed the influence of privacy awareness. Through comparing the behaviour of the public against our researcher sample, we determined how actions vary. The other four directly related to our first research subquestion: *Can the Privacy Paradox be confirmed in the UK?* To explore whether it exists, each hypothesis compares claims against behaviour. We triangulate findings by assessing concerns, protection and disclosure. If a disparity is found consistently, it is likely that the Paradox is present.

Hypotheses: Behavioural comparisons. We begin by discussing our comparisons between the two participant groups. For our first hypothesis, we believed our expert sample would protect themselves more often than the general public. This expectation emerged for two reasons. To use privacy-protecting tools, individuals should perceive a degree of risk. They must also understand the purpose of these tools and how they are operated. We thought security experts were more likely to fulfil both of these criteria. Therefore, we asserted that the researcher sample will have a significantly higher mean *behaviour score* than the public sample. This supports our Paradox analysis by examining a potential contributory factor.

Secondly, we expected the researchers to disclose less information than the general public. These individuals should understand the risks to user privacy. They might also have knowledge of data sharing and inference techniques. Therefore, they should appreciate the importance of cautious behaviour. Based on this rationale, we posited that the researcher sample will have a significantly lower mean *disclosure quantity* than the public sample. This further examines a contributory factor, supporting our central research question.

Hypotheses: Paradox confirmation. We now move on to consider the relationship between claims and participant behaviour. These analyses concerned the public group, since the expert group had a small sample size ($n = 25$). For this size, significant correlation would be unlikely even if an association existed. Furthermore, we sought to confirm the Paradox in the public, rather than for researchers. Therefore, our final four hypotheses concern the street sample.

We did not expect any association between self-perceptions and online protection. Even if an individual claims to act privately, this might be biased by several factors. They might be prone to optimism bias and overestimate their protection [220]. They might lack privacy awareness, and therefore be unaware of settings. For these reasons, we asserted that there would be no significant correlation between *perception scores* and *behaviour scores*. A similar correlation approach was used by Zafeiropoulou et al. [397] in their analyses. We accept that a lack of significance does not imply an absence of relationship. Therefore, we also check that any correlation coefficient r_s is less than ± 0.1 . This would imply a ‘small’ effect size [94], and that any relationship is very weak [366]. This hypothesis had direct relevance to our research questions, as it compares claims against actions. If accepted, a disparity might be present.

We also doubted that these perceptions would correspond with empirical actions. Our optional questions gave participants an opportunity avoid disclosure. The fields were explicitly optional, and a ‘Prefer not to say’ field provided a simple alternative. Intuitively, if a person claims to act privately, one would expect them to minimise their disclosure. However, the Privacy Paradox would contradict this rational assumption. To explore the matter, we posited that there would be no significant correlation between *perception scores* and *disclosure quantity*. We also sought to receive a $r_s < \pm 0.1$. If we accept this hypothesis, we further demonstrate that a disparity is present.

We believed that concerns would bear no relationship with online behaviour. Surveys have shown that individuals claim the importance of privacy [304]. However, even when concerns are expressed, users often lack protection [4]. Therefore, we

asserted that no significant correlation would exist between *concern scores* and *behaviour scores*. Again, we expected a $r_s < \pm 0.1$. If evidence supports our conjecture, credence is given to the Paradox.

Finally, we doubted that concern would influence empirical behaviour. Previous work demonstrates that individuals tend to reveal far more than they claim [279]. Therefore, we posited that there would be no significant correlation between *concern scores* and *disclosure quantity*. We again anticipated a $r_s < \pm 0.1$. This comparison of claim and behaviour sought to further demonstrate a disparity.

Our six hypotheses are formally summarised below:

1. The mean *behaviour score* of the researcher sample will be significantly higher than the mean *behaviour score* of the public sample.
2. The mean *disclosure quantity* of the researcher sample will be significantly lower than the mean *disclosure quantity* of the public sample.
3. For the public sample, there will be no significant correlation between the *evaluation scores* and the *behaviour scores*.
4. For the public sample, there will be no significant correlation between the *evaluation scores* and the *disclosure quantity*.
5. For the public sample, there will be no significant correlation between the *concern scores* and the *behaviour scores*.
6. For the public sample, there will be no significant correlation between the *concern scores* and the *disclosure quantity*.

All six of our hypotheses directly further the central research question. If several of these conjectures are confirmed, it would suggest a disparity between claims and behaviour. This would confirm the existence of the Paradox, enabling us to move forward to explore the IoT. If the expert group acts in a cautious manner, perhaps education could encourage protection.

Survey design

In this subsection, we discuss the questions asked in our survey. The questionnaire was composed of five sections: *Perceptions*, *Online behaviour*, *Mandatory demographics*, *Optional demographics* and *Informed consent*. These queries and their sequence can

Table 4.1: Street Survey Questionnaire

#	Mandatory demographics					
1	What is your gender?					
	<i>Male</i>	<i>Female</i>	<i>Other</i>			
2	What is your age range?					
	<i>18-25</i>	<i>26-35</i>	<i>36-45</i>	<i>46-55</i>	<i>56-65</i>	<i>66 and over</i>
3	What is your highest education?					
	<i>School</i>	<i>GCSE</i>	<i>A-Level/College</i>	<i>Degree</i>	<i>Masters</i>	<i>PhD</i>
4	What is your country of residence?					
	<i>(Open Reply)</i>					
#	Optional demographics					
1	What is your marital status?					
	<i>Single</i>	<i>Married</i>	<i>Divorced</i>	<i>Widowed</i>	<i>Prefer Not To Say</i>	
2	What is your Twitter username?					
	<i>(Open Reply)</i>		<i>Prefer Not To Say</i>			
3	Who is your employer?					
	<i>(Open Reply)</i>		<i>Prefer Not To Say</i>			
#	Perceptions(Strongly Agree to Strongly Disagree via Neutral)					
1	Online privacy is of importance to me.					
2	I am private with my data.					
#	Online behaviour (Always, Often, Rarely, Never, Unsure, N/A)					
1	How often do you clean your Internet browser’s history?					
2	How often do you use Internet browser plug-ins/extensions to protect your privacy?					
3	How often do you encrypt data on your computer?					
4	How often do you store unencrypted data within a cloud provider such as Dropbox?					
5	How often do you share public posts on social networking sites such as Facebook?					
6	How often do you share your location on social networking sites?					
7	How often do you use Tor for your web browsing?					
8	How often do you use encryption tools for your emails?					
9	How often do you read the terms and conditions on websites you use?					
10	How often do you check permissions before installing smartphone apps?					

be found below in Table 4.1. We now continue by outlining the rationale which informed their selection.

Section order. The sequence of questions was important, as order effects can influence responses [360]. When the topic of privacy is primed, individuals often report greater apprehension [65]. Since we sought to analyse baseline concerns, we did not want reactions to be magnified. Therefore, before mentioning ‘privacy’ in a question, we assessed disclosure behaviour in the *Optional demographics*. When it was then included in the *Perceptions* section, it was relevant to these queries. *Online*

behaviour was discussed later, after participants had already expressed their concerns. As responses were made with pen on paper, individuals could not return to adjust their opinions. This approach should mitigate potential order effects.

Perceptions. We used our *Perceptions* section to analyse the views of each participant. Through assessing agreement with simple statements, we explored both privacy concerns and self-perceptions. To gauge concern, we asked participants whether they considered online privacy to be important. Although we evaluate concern, we made effort to avoid this terminology. This was crucial, as we wished to disguise the true focus of our study. Otherwise, demand characteristics could lead participants to answer in a favourable manner [287]. Fortunately, if a person deems this issue to be important, then they express concern for the matter. As previously mentioned, responses were made on a five-point scale from *Strongly Agree* to *Strongly Disagree*.

To solicit self-perceptions, we asked each individual whether they were private with their data. The term ‘private’ was preferred to ‘privacy-protective’, as the latter was considered excessively complex. Again, agreement was ranked on a five-point Likert Scale. Concerns and self-perceptions were assessed through a single question each. This approach sought to minimise questionnaire length, as this can contribute to response fatigue [88]. We deemed this important in street surveys, since interactions are limited on time. While this was appropriate in this confirmatory study, concerns are deconstructed in later chapters.

Online behaviour. For behaviour, we explored the usage frequency of protective tools. The chosen techniques were varied, ranging from encryption to anonymous browsing to app permissions. We did not adopt the selection criteria from Section 3.5. In this initial study, we were not comparing behaviour across devices. Furthermore, our focus was not scoped to a subset of devices. We simply explore high-level actions to confirm the Paradox. Therefore, we selected tools through three basic criteria: *utility*, *comprehensibility* and *commonality*. Firstly, it was crucial that each technique was beneficial for privacy. If it did not offer information or protection, then it was not relevant to our considerations. Secondly, the approach should be comprehensible to our audience. To support this, we avoided the use of technical terminology. Although anonymous browsing requires some knowledge, most of the settings were common. Finally, the tools should be frequently found in digital environments. They might be included on existing software (e.g., web browsers) or compatible with many operating systems. This grants individuals an opportunity to use these features. Through these criteria, we selected the above seven privacy-protective techniques (all *Online behaviour* questions other than 4, 5 and 6).

We had concerns that participants might blindly respond ‘Always’ for all the techniques. This assimilation effect would lead to inaccurate judgements of behaviour [342]. To mitigate bias, we also included actions which place privacy at risk. To disguise this approach, the question order was shuffled. The actions are found in *Online behaviour* questions 4, 5 and 6. These activities were also selected based on the three aforementioned criteria. Since ‘Never’ responses were now privacy-protective, these replies were reverse-coded. Through assessing the frequency of varied actions, we conducted a broad analysis of participant behaviour.

As previously mentioned, tool usage was reported on an ordinal scale (spanning from ‘Always’ to ‘Never’). Therefore, rather than assessing whether techniques were used, we also considered the frequency of usage. This accommodated the fact that settings could be accessed once but quickly forgotten. We also included ‘Unsure’ and ‘N/A’ options to avoid compelling uninformed replies. As it would be inappropriate to punish uncertainty, such responses were not factored into *behaviour scores*. When calculating these scores, we took a mean of the frequencies (with reverse-coding applied). This resulted in a figure from 1 (low) to 4 (high). If an individual took frequent action (4/4), then they protected their privacy. In contrast, if tools are rarely used (1/4), data might have been placed at risk.

Mandatory demographics. Demographic fields were divided into two sections: *Mandatory demographics* and *Optional demographics*. We used the former to monitor the representativeness of our study. While we also undertook demographic comparisons, all discussion is omitted from the thesis for brevity.

Our form began by soliciting gender, age range, highest education level and country of residence. None of these queries were considered to be particularly sensitive, and all are commonly found in questionnaires. This was important, as it avoided priming concerns before the optional section.

Optional demographics. These queries were included to assess empirical disclosure behaviour. Three details could be supplied, with fields displayed in order of sensitivity. Each field was optional, with a ‘Prefer not to say’ field offering a simple alternative. We also informed the participants of this situation through an oral script. Therefore, if a person chooses to disclose data, they do so of their own accord. We accept that there is a fabrication risk, with individuals feigning cooperation. However, since it is quicker and easier to avoid the question, we doubt false details were frequently given. This matter is discussed in the critique (Section 9.3).

The first field was marital status, regarded under the Data Protection Act and GDPR as ‘personal data’¹. This was chosen as it reveals information about the participants’ private life. As this query was optional, participants would be wise to avoid unnecessary disclosure. The second field was for Twitter username; data which could be used to research the individual [11]. This was included as it reveals details about a participant’s online identity. Since the vast majority of Twitter accounts are public [255], this could enable access to personal information. These profiles can host large quantities of data, including images, websites and private opinions. This could allow sensitive details to be discovered or inferred. The username could also be used to contact the individual, which we would expect to be undesirable. Finally, we solicited the identity of the individual’s employer. This information reveals details about a person’s professional life. Such data, particularly in combination with gender and age range, could make an individual highly identifiable. It could also be used to infer salary range, office location and socio-economic status. Since avoiding disclosure was simple, we did not expect this to be revealed.

The number of revealed details denoted each individual’s *disclosure quantity*. Importantly, we assessed the presence or absence of this data, rather than the data itself. This approach was taken for ethical reasons, as this information could be sensitive. Although disclosure is not a perfect proxy for behaviour, it does give empirical evidence of a person’s actions. If a person claims to value their privacy, it would be rational for them to show caution.

Informed consent. We then solicited consent from our participants. We began by explaining our data collection and how responses will be stored. We continued by highlighting redressal procedures and how to withdraw from the study. To minimise confusion, an oral script was also delivered by the researcher. Finally, both parties signed the document to indicate their agreement. If an individual forgot to complete this section (which happened on three occasions), their responses were securely discarded. Therefore, we only analysed data from those who gave informed consent. Before undertaking the study, we received ethical approval from our institution. We also got explicit authorisation from the four local councils for each site. By complying with these requirements, we undertook our research in a responsible manner.

¹<https://www.policybee.co.uk/blog/13346/new-data-protection-regulation-out-with-the-data-protection-act-in-with-gdpr>

Participant recruitment

Public sample. We recruited participants across four distinct sites: London, Birmingham, Cardiff and Oxford. Since we wished to survey members of the UK public, this was preferred to targeting one location.

To solicit a heterogeneous sample of citizens, we selected sites varying in region, population size and ethnicity. London has a population of over 8 million people and a diverse mix of ethnicities². As the largest metropolitan area in Western Europe, it typifies an urban environment. Our site was based on a busy shopping street in the borough of Hammersmith & Fulham. Birmingham has just over 1 million inhabitants and is the most-diverse city in the UK³. Socio-economic status tends to differ from London and it retains a strong Midlands identity. We distributed questionnaires near Birmingham New Street Train Station, since it is the busiest station outside of London. Cardiff has 341,000 people and is predominantly White British⁴. It is also in Wales (rather than England), with the nation possessing a distinct culture. The survey took place on St Mary's Street, the central commercial street of the capital. Finally, Oxford is least metropolitan, with a population of 150,000⁵. Our site was situated on Cornmarket Street, a busy pedestrianised thoroughfare.

For logistical reasons, no sites were included from Scotland or Northern Ireland. In future work, we wish to target larger samples in a diverse range of locations.

The doctoral researcher recruited all the participants, mitigating influence from multiple individuals. They also distributed the forms and acquired informed consent. Questionnaire distribution took place for five hours daily between 10am and 3pm. These periods were chosen since: a) the sites were most busy; and b) it allowed for travel time. Since the participants were recruited on public streets, we directly interacted with a varied demographic.

Our screening criteria were simple: all participants must be adults who are able to give consent. We decided against a restricted sample since we wished to survey the general public. Individuals were not allowed to participate more than once, as that would bias the distribution of our results. On the rare occasions that respondents forgot to sign their form, their responses were discarded.

The public were compensated with a nut-free chocolate bar for their time. We felt this was appropriate since form completion took approximately three minutes. Al-

²<https://www.trustforlondon.org.uk/data/topics/population-geography/>

³<http://worldpopulationreview.com/world-cities/birmingham-population/>

⁴<http://worldpopulationreview.com/world-cities/cardiff-population/>

⁵https://www.oxford.gov.uk/info/20131/population/459/oxfords_population

though this approach incentivised participation, we doubt that responses were unfairly influenced. Each bar only cost 30p and individuals were free to reject the compensation. Indeed, many refused the reward and were more interesting in providing their opinion. Since the purpose of our study was disguised, demand characteristics should have little influence [287].

Researcher sample. In this initial chapter, we sought to confirm the existence of the Privacy Paradox. However, we also took an analysis of contributory factors. As previously mentioned, a lack of awareness might contribute to the issue [48,309]. We would therefore expect those with privacy knowledge to act in a more-protective manner. This is because they should recognise the risk and know the response. To explore this, we must solicit the opinions of privacy experts.

In addition to our public sample, we recruited a group of cyber security researchers. These individuals were contacted via email and invited to complete a physical questionnaire. They were local to the student's university and did not receive compensation. The chocolate bar was deemed unnecessary since the individuals already had an interest in academic research. However, to mitigate the influence of demand characteristics [287], the focus of the study was not revealed. These participants all possessed good privacy awareness through their education and research. As both samples were composed of UK residents, comparisons should not be biased by geographic culture.

4.3 Results and Discussion

Participants

Public sample. Across the four sites, we surveyed a total of 112 participants. Three additional responses were discarded due to questionnaires not being signed. 105 (93.8%) of the 112 identified themselves as UK residents. Of the remaining seven, two came from Australia, two from Ireland, and one each from Spain, India and the United Arab Emirates (UAE). We found no significant differences between the responses of UK residents and those from other nations. Since the vast majority resided in this country, we interacted with members of the UK public.

The gender ratio was quite balanced, with 57.1% identifying as female and 42.9% identifying as male. The most popular age ranges were 18-25 (25.9%) and 66 and above (24.1%). At least 9% of individuals were in each group, with an estimated mean age of 43.3. Since the UK median is 40 years old [190], our sample appears to bear some relevance.

Other than PhD, at least 10% of our participants came from each educational level. 37.5% received an undergraduate degree, while 25.9% finished after secondary school. With a further 13% possessing a Master's, this sample is relatively well-educated.

Researcher sample. We recruited 25 cyber security researchers for our expert sample. The male-female split was 80%-20%, which is somewhat representative of the current industry⁶. 80% of participants were under 36 years of age, likely because we recruited from a university environment. This contributed to an estimated mean age of only 31.2. Being academics, all individuals had at least a Bachelor's qualification. 76% possessed a Master's degree and 8% had completed a doctorate. As previously mentioned, all these participants were UK residents. This supported comparisons of opinion with our public sample.

Privacy perceptions

Public sample. We now move forward to discuss the concerns of the public sample. Of our 112 respondents, 104 (92.8%) agreed with the importance of online privacy. This was compared to only 3 (2.7%) who disagreed with the statement. Furthermore, since 80 expressed strong agreement, 71.4% were particularly concerned about their privacy. When calculating the *concern scores*, participants received a mean of 4.62/5. We then used a single-sample Wilcoxon Signed-Rank Test to compare our distributions to the neutral median (3/5). Since we received $p < 0.001$, our results are unlikely to be due to chance. Through our hypotheses, we shall explore whether behaviour is commensurate with these concerns.

When considering self-perceptions, the public appeared more reserved. However, 91 (81.3%) still claimed to act privately with their data. 40.2% expressed strong agreement, while only 8% showed disagreement. When calculating the *perception scores*, the mean was 4.13/5. This suggests that most of our sample claim to behave in a private manner. Through another single-sample Wilcoxon Signed-Rank test, these distributions differed significantly from the neutral median ($p < 0.001$). Therefore, we can surmise that our public sample claim to act privately. While one might expect corresponding behaviour, we posit that protection will be rare.

Researcher sample. We now consider the views of privacy experts. We expected their concerns to also be strong, since it is a social norm to value the principle [352]. When analysing the responses, 92% of researchers agreed with privacy's importance. 52% expressed strong agreement, with this lower than the proportion in the public

⁶<http://www.govtech.com/workforce/Why-Are-So-Few-Women-in-Cybersecurity.html>

sample (71.4%). No respondents disagreed with the statement, contributing to a mean *concern score* of 4.44/5. Again, this metric is slightly lower than that of the general public (4.62/5). Although this difference was not significant, experts' concerns might be more nuanced.

We next analysed self-perceptions, assessing whether the researchers felt they acted privately. We expected lower scores than for the public sample, since experts should have greater self-awareness. They might also be less prone to optimism bias [29], as they should understand the risks that they face. When asked if they acted privately, 80% agreed with the statement. However, only 28% expressed strong agreement, compared to 40.2% in the public sample. 12% disagreed that they were private, again responding more-negatively than the general public (8.0%). While these differences were not significant, privacy awareness might contribute to balanced judgements. The *perception score* was 4.2/5, suggesting most respondents considered themselves private. We would expect them to protect their data, and this is explored in the following section.

Online behaviour

Public sample. We continue by discussing the online behaviour of our public sample. Less than 50% of participants cleaned their browsing history (at least) often, compared to 23.2% who always acted. This suggests that most individuals erase their history rarely. Use of browser extensions was even less frequent, with only 36.6% using them often. This is unfortunate, as tools such as NoScript can limit online tracking. While encrypting data is valuable for protection, 66.1% never performed this task. This was compared to only 6.3% who always act, indicating an absence of encryption. We found this inaction surprising, especially considering the efforts of public awareness campaigns.

When discussing more-complex tools, action was also rare. 75.9% of individuals had never used Tor, compared to only 4.5% who use it often. Email encryption was just as infrequent, with 72.3% never taking this step. These results are unsurprising since the tools require additional knowledge. However, we thought that Tor media attention might have generated interest [50, 77]. As both applications have been criticised for usability [91, 388], their interfaces might impede adoption.

Individuals have two common opportunities to increase their awareness: privacy policies and app permissions. They both support knowledge of the threat and use of protection. Despite this, 40.2% admitted to never reading the Terms and Conditions on websites. A further 22.3% checked them rarely, with only 14.3% doing this always.

This is supported by a 2017 poll, which found 96% tend to avoid these policies [155]. If true, millions of individuals are held to conditions of which they have little knowledge. Similarly, 30.4% never check smartphone permissions, compared to 19.6% who always do. As these settings outline data access, almost a third of our sample remain oblivious.

When considering ‘risky’ practices, we applied reverse coding. Encouragingly, only 33.9% have uploaded unencrypted data to the cloud. However, the 17.9% who always do might be placing their information at risk. Since cloud services were less common at the time of the survey (August 2015), this figure might have since increased [169]. Only 36.6% shared public posts on Facebook, with 25.9% never undertaking this action. This suggests that many individuals refrain from this active exposure. Participants were even more protective over their location data. Only 14.3% shared it online, reducing the risks from identification and tracking.

Our results present a juxtaposition between taking action and disclosing data. Whereas few participants actively shared their details, protective tools were inconsistently used. This inertia might be due to apathy or default settings [244]. It might also derive from a lack of awareness, with users having insufficient skills. This can be partially explored through comparisons with the expert sample.

We found distributions differed significantly from the median ($p < 0.001$), apart from for web browsing, Facebook sharing and phone permissions. This suggests that our results were unlikely due to chance. For the public sample, the mean *behaviour score* was 2.16/4. While respondents used some protection, they could do far more to guard their data.

Researcher sample. Since the researchers had greater privacy knowledge, we expected them to appreciate protection. They were indeed found to use many tools more frequently. Browser extensions were used often by 60%, with 40% enabling them constantly. As only 15.2% of the public acted similarly, the frequency differed significantly ($U = 803$, $p = 0.001$, $d = 0.593$). Researchers were also more likely to encrypt their files ($U = 475$, $p < 0.001$, $d = 0.981$), since 60% did this often. This might be because they possessed awareness of the threats they face.

There was a large contrast in the usage of advanced tools. 12% of experts used Tor often, which was still significantly more frequently than the public ($U = 766.5$, $p < 0.001$, $d = 0.633$). Finally, researchers made greater effort to read app permissions. 52% undertook this action consistently, compared to only 19.6% in the public sample. This significant difference ($U = 701.5$, $p < 0.001$, $d = 0.705$) suggests that the experts are better informed on smartphone practices.

H1. Based on the protective actions of the researchers, we calculated their *behaviour score*. This group received 2.52/4, implying that they take some steps to protect their data. While their actions were not consistent, some tools were used frequently. We now turn to our first hypothesis, which asserts that the researcher sample will have a significantly larger mean *behaviour score* than our public sample. If accepted, privacy knowledge might support protective behaviour. A Mann-Whitney U Test found that the researchers took action significantly more frequently ($U = 672.5$, $p < 0.001$, $d = 0.738$). Therefore, we accepted this hypothesis. The ‘large’ effect size [333] further emphasises the contrast in behaviour. As discussed above, concern did not vary between our two samples. Since the Paradox describes the concern-behaviour disparity, it might be smaller in the expert group. If privacy awareness is influential, education might provide a mitigative approach.

Empirical disclosure behaviour

We have considered the online actions of both our samples. However, to empirically explore behaviour, we analysed the *disclosure quantity*. As previously mentioned, all three of the fields were explicitly optional. This was highlighted both on the form and through an oral script. Immediately after these fields, over 80% of respondents claimed to act privately. Therefore, one might expect scarce disclosure. However, we posit that a disparity exists between these claims and behaviour.

Public sample. In terms of public disclosure, 96.4% revealed their marital status. This was compared to only 3.6% who opted to omit the field. This data reveals information about the participants’ private life, and could be used in combination with other details. Therefore, it might be wise to withhold this information. The next field concerned Twitter username, with a ‘Prefer not to say’ option clearly provided. Although only 11.6% disclosed this data, only 12.5% explicitly chose not to. The remaining 75.9% lacked an account, possibly due to the older demographic of the group. Since the study was conducted (August 2015), Twitter’s user base has increased⁷. Therefore, disclosure might be greater if the survey was repeated. Those revealing their username enabled direct contact, as most accounts are publicly-accessible [255]. Through details on this portal, such as images and personal opinions, much could be learned about an individual. We accept that account absence could be falsely claimed to avoid disclosure. However, since it was quicker and easier to ignore the field, we doubt this was common.

⁷www.statista.com/statistics/282087/number-of-monthly-active-twitter-users/

The final field considered a participant’s employer. This reveals details concerning their professional life, and could be used to infer information. Despite this, 83.9% disclosed their company or work status. Since only 16.1% chose to conceal this data, the sample appeared willing to share. By combing work details with other demographics, further details could be discovered. For example, if a person’s company, age range and gender is known, one could filter their identity to a small subset [106]. Social networking sites, such as LinkedIn, could also host additional information. Despite all three fields revealing personal details, participants chose to disclose their data.

We conducted single-sample Chi-Squared tests to analyse whether the proportions were due to chance. Since we received $p < 0.001$ in all three cases, our distribution differed significantly from an even division. The modal *disclosure quantity* was 2/3 items, with a mean average of 1.93. Although most of the public claimed concern for privacy, they avoided opportunities to protect their data. Over 99% disclosed a personal detail needlessly, with 72.3% revealing two elements. As most participants claimed to act in a private manner, this might suggest claims differ from behaviour.

Researcher sample. We next explored the empirical behaviour of the researchers. Disclosure was expected, since it is a social norm to share information [127]. However, as this group had privacy expertise, they might show greater caution. Similar to the public, 96% disclosed their marital status. This suggests that this information was not valued highly. When it came to Twitter usernames, behaviour differed greatly. Although experts revealed details more-frequently (28%), a large proportion explicitly refused (64%). Only 8% claimed to lack an account, and membership might be expected of the young and educated demographic⁸. Since rejection was 5x greater in this sample, researchers might be more cautious with their accounts. Finally, only 56% disclosed their employer, compared to 83.9% in the public group. This is surprising, as their academic affiliation might not reveal sensitive data. The contrast in proportions suggests that researchers are more private with their information.

H2. For our second hypothesis, we posited that this researcher sample will have a significantly lower mean *disclosure quantity* than the public sample. This was expected since those who recognise risks should act more cautiously. For our academic group, the modal *disclosure quantity* was again 2/3 items. When comparing the groups through an independent-samples t-test, we found no significant difference ($t(27.989) = 0.677, p = 0.504$). Therefore, we failed to accept this hypothesis. We attribute this similarity to Twitter disclosure, since a greater proportion of researchers used the service. Furthermore, our small sample size ($n = 25$) impeded the likelihood

⁸<https://blog.hootsuite.com/twitter-demographics/>

of significance. Despite this result, experts still take greater action online to guard their data. This might imply that awareness can encourage protective behaviour.

The Privacy Paradox

We now seek to address our first research subquestion: *Can the Privacy Paradox be confirmed in the UK?* This is achieved by testing this chapter's final four hypotheses. These questions all analyse the disparity in a different manner, comparing *privacy concerns* and *self-perceptions* with *online behaviour* and *empirical behaviour*. If a disparity appears consistently, we can confirm the existence of the Paradox. As previously mentioned, these analyses concern the public sample rather than the researchers. As our research seeks to support users in protecting their data, the non-expert group has greater relevance.

H3. Our third hypothesis asserted that no significant correlation would be found between the *perception scores* and the *behaviour scores*. If accepted, this would imply that claims of privacy have no bearing on protective behaviour. Through Spearman's Rank technique, we found no significant correlation between these two factors ($r_s = 0.078$, $p = 0.423$). Therefore, we accepted this hypothesis, with it suggesting a disconnect between claims and actions. As we received a $r_s < \pm 0.1$, any effect would have been minor regardless. To illustrate the disparity between the factors, we produced the below heatmap (Figure 4.1). Based on the distribution of responses, we can view that most individuals claim to act privately. This is represented by the darker red shading near the top of the graphic. However, most of these participants also received low *behaviour scores*. Indeed, over half of the group used protection rarely despite their strong claims. With the top-left region being extensively populated, this highlights the issue.

H4. For our fourth hypothesis, we posited that there would be no significant correlation between the *perception scores* and the *disclosure quantity*. If an individual claims to act privately, one would expect them to avoid exposure. However, if such claims are made seconds after needless disclosure, a disparity might exist. 68.8% reported acting privately but revealed two elements of data. When considering any elements, 80.4% acted in this manner. Through another Spearman's Rank test, we found no significant correlation between *perception scores* and *disclosure quantity* ($r_s = -0.062$, $p = 0.518$). Therefore, we accepted our fourth hypothesis. Since $r_s < \pm 0.1$, any effect would have been negligible. We illustrated the disparity through another heatmap, found below in Figure 4.2. The shaded top presents how most respondents claimed to act privately. However, the rightward distribution shows how




Figure 4.1: Heatmap: Self-Perceptions vs Online Behaviour

many disclosed unnecessarily. The bottom-left region is sparsely populated, with most responses grouped in the top-right. This disparity might further suggest the existence of the Paradox.

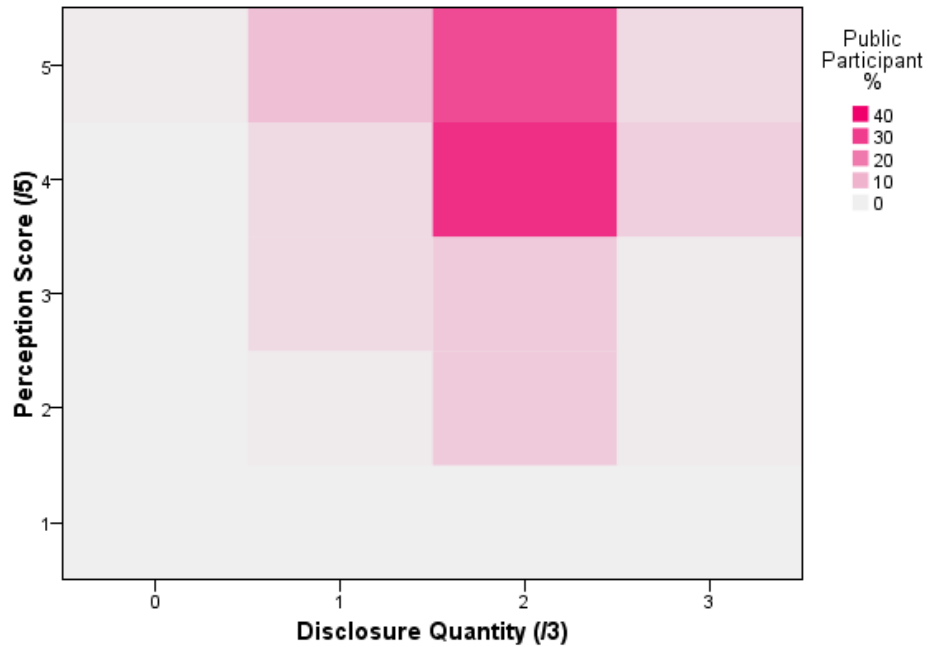


Figure 4.2: Heatmap: Self-Perceptions vs Empirical Behaviour

H5. We then moved forward to consider privacy concerns. Our fifth hypothesis asserted that no significant correlation would exist between *concern scores* and *behaviour scores*. If true, this would imply that stated opinions have little relationship with protection. This counter-intuitive notion would support the Paradox. A Spearman’s Rank test found that there was no significant correlation between the two factors ($r_s = -0.010$, $p = 0.916$). Therefore, we also accepted this hypothesis. As $r_s < \pm 0.1$, significance does not appear solely impeded by sample size. To illustrate the contrast, we produced another heatmap (Figure 4.3). The topmost region is heavily shaded, displaying that participants express strong concerns. However, protection was inconsistent, hence skewing shades to the top-left. As the top-right is sparsely populated, a concern-behaviour disparity appears present.

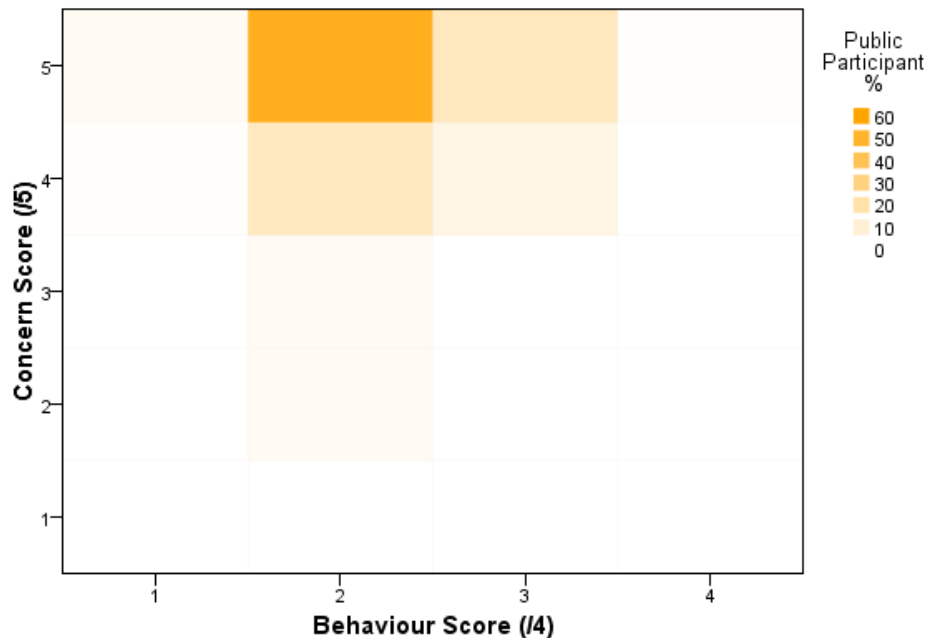


Figure 4.3: Heatmap: Privacy Concern vs Online Behaviour

H6. Finally, we compared privacy concerns against empirical behaviour. This explores the key disparity underlying the Privacy Paradox. Our sixth hypothesis posited that no significant correlation would be found between *concern scores* and *disclosure quantity*. Supporting this, 75.8% of our respondents both claimed concern and revealed two elements of data. This was compared to only 17.0% who acted in a cautious manner. Although the fields were not critically sensitive, disclosure was strictly optional. As before, we found no significant correlation between the aforementioned factors ($r_s = -0.146$, $p = 0.124$). Therefore, we accepted this hypothesis along with the other Paradox conjectures. Although $r_s \not\leq \pm 0.1$, any influence was

probabilistically due to chance. It appears that concerns have little relationship with empirical behaviour. We produced a final heatmap, with this graphic shown below in Figure 4.4. The shaded region illustrates how our respondents expressed strong concern. However, the distribution is rightward shifted due to frequent disclosure. As the top-right is heavily populated, this suggests the presence of the Privacy Paradox.




Figure 4.4: Heatmap: Privacy Concern vs Empirical Behaviour

Hypotheses H3-H6 explored our first subquestion (*Can the Privacy Paradox be confirmed in the UK?*) from several angles. Since these conjectures have been accepted, it appears that concerns bear little relationship with behaviour. We do not claim to conduct a representative study of the UK public. However, our findings suggest that many residents are susceptible to the disparity. Therefore, the Privacy Paradox has been confirmed. We can now move forward to explore the IoT.

4.4 Implications

Paradox. To explore the Paradox in the UK, we conducted a 112-person street survey. 93% of our sample valued privacy, suggesting a strong degree of concern. Over 80% reported acting privately, also implying the principle is appreciated. However, these claims do not seem to align with behaviour. When considering online activity, individuals fail to protect themselves consistently. Permissions are rarely checked and encryption is infrequently used. This inaction might place personal data at risk.

Empirical behaviour also appeared incommensurate with claimed concerns. Over 99% disclosed data needlessly, with 76% revealing more than one element. These disclosures were made seconds before individuals reported acting privately. In our four final hypotheses (H3-H6), behaviour consistently differed from claims. This supports the presence of the Paradox.

Awareness. We also recruited a sample of cyber security researchers. Since they had expertise, we explored the possible influence of privacy awareness. These individuals were found to have strong concerns and self-perceptions. Furthermore, their behaviour was more-protective in several regards. Researchers encrypted their files more frequently ($p < 0.001$) and read app permissions more often ($p < 0.001$). While disclosure was similar to the public sample, experts showed greater caution with Twitter usernames and employment details. Since privacy literacy can encourage protection [48], increasing awareness might be an effective approach.

Next steps. This chapter concerned technology in general, rather than the Internet-of-Things. Now that the Paradox has been confirmed, we can continue by exploring IoT perceptions. Furthermore, since privacy awareness might aid protection, we will consider this in later studies. Importantly, settings appear to be inconsistently used. To help users reduce their risk, we should support protective behaviour.

Smart devices present a variety of risks [126, 181, 300], with these issues well-documented in the media [131, 154, 273]. A ‘fear of the unknown’ might also heighten consumers’ concerns [221]. This could increase the first component of the Paradox. Due to the novelty of the IoT, it remains relatively unfamiliar [42]. The usability of such devices has also been criticised on several occasions [90, 143, 183, 380]. This could constrain protective behaviour [317], impacting the second Paradox component. As the disparity is confirmed in this chapter, we can now explore the IoT.

To examine the idiosyncrasies of smart devices, we require a comparative baseline. Therefore, we also solicit the opinions towards less-novel computers. Privacy concerns tend to be strong, as highlighted in this chapter. If the IoT presents behavioural constraints, the Paradox might be greater in this environment.

Concerns can be gauged through questionnaire responses. However, behaviour is more challenging to assess. Functionality differs across the spectrum of devices, constraining analyses of protection. Therefore, we seek to explore usability and familiarity, since they can be influential [158, 230, 317]. The act of buying a product might also be indicative of behaviour. Through assessing purchasing decisions, we can explore whether privacy influences choices. To analyse IoT perceptions, we survey a large range of devices. This provides an initial insight into this nascent environment.

Chapter 5

How do perceptions differ between IoT and other devices?

5.1 Introduction

Background

Through our initial study, we sought to confirm the Paradox in the UK. Since we have addressed our first research subquestion, we proceed to analyse the Internet-of-Things. Research has suggested that the IoT might challenge privacy-protective behaviour. Firstly, studies have shown that smart devices often suffer from usability issues [143]. When tools are excessively complex, costly mistakes can easily occur [158, 230]. Secondly, due to the novelty of the environment, IoT devices are quite unfamiliar [42]. Products might differ from established mental models, particularly when the range is heterogeneous. While consumers may have learned familiar settings, the IoT poses a new challenge. Thirdly, smart devices are frequently driven by data collection [300]. Wearables and home automation deliver much of their functionality through their sensors. As devices ubiquitously inhabit our homes, the opportunities for surveillance will only increase [108].

We expect privacy concerns will be strong for several reasons. Firstly, the risks of the IoT have been documented for several years [131, 154, 273]. Secondly, unease might naturally arise due to a ‘fear of the unknown’ [221]. Finally, individuals will continue to aspire to privacy due to social norms [352]. If behaviour is constrained and concerns develop, the Paradox will be exacerbated. This might lead users to inadvertently place their data at risk [194].

Motivation

At this stage we have confirmed that the Paradox exists in the UK. However, we have not examined the influence of the IoT. This matter is directly concerned by our second research subquestion: *How do perceptions differ between IoT and other devices?* To address this query, we require a baseline for comparisons. Therefore, we explore a range of products through an online survey. Once high-level differences are understood, we can dissect rationale through qualitative interviews.

Surveys were deemed most appropriate for this broad exploration [130]. Again, we targeted members of the UK public. Our street survey (Chapter 4) was valuable for accessing a diverse demographic. However, the sample size was limited and respondents were particular to certain cities. To survey a larger group, and across the nation, we selected an online approach. At this stage, we did not require participants to be device owners. Such criteria would have constrained sample size and impeded a broad evaluation. We simply wished to compare general perceptions.

When evaluating the Paradox, concerns and behaviour are considered. In this high-level survey, concerns can be gauged through Likert Scale ratings. By requesting privacy evaluations, low rankings highlight the greatest worries. This approach was preferred to ‘concern ratings’, which might be inflated through social desirability bias [65]. Behaviour was more challenging to analyse, due to the wide diversity of devices. Therefore, we considered factors that might influence behaviour. Both usability and familiarity could also be assessed through ratings. If the IoT scores poorly, these issues might impede protection [158,230,317]. We also explored whether individuals decided to purchase the product. This is a reasonable proxy, since behavioural rationale is used to undertake this action. If a person doubts a device’s protection, one might expect them to reject it. However, if they express concerns but still acquire it, then claims might differ from behaviour.

Purchases can provide some indication of participant behaviour. However, their rationale is of even greater interest. Justifications can differ greatly and highlight how decisions are made. Therefore, unlike the previous chapter, we also collect qualitative data. This is achieved by soliciting reasons for product ownership (or lack thereof). If purchases are driven by a desire for functionality, protection might be an afterthought. We can also identify whether rejections are due to privacy or other considerations. If privacy is frequently disregarded, this might help explain the Paradox.

5.2 Methodology

Hypotheses

Factors. We begin by discussing our four hypotheses and highlighting how they relate to our current subquestion: *How do perceptions differ between IoT and other devices?* The compared products were: *desktops, laptops, tablets, wearables, smart appliances* and *home automation*. Their selection is justified later in this section. However, we first outline the factors that we analyse.

In our online survey, we wished to explore perceptions towards a range of devices. However, if privacy is primed, concerns might be inflated [65]. To mitigate this issue, we sought to disguise our study focus. Fortunately, in addition to *privacy*, we explored three additional factors. These comprised *usability, familiarity* and *utility*. The study was also framed as concerning technology in general, rather than focusing on privacy. This sought to minimise bias from demand characteristics [287].

Previous work suggests that complexity can constrain privacy-protective behaviour [388]. If a product is challenging to use, a person might reveal data through mistakes or misconfigurations. Protective settings may also be avoided if the interface is not easy to understand [59]. Therefore, we wished to explore perceptions of *usability*. If the IoT is regarded as challenging, protection might be limited.

Unfamiliar devices can also cause privacy issues, as they might fail to align with mental models [317]. Moreover, when a product is not understood, settings might be avoided. Inexperience might also lead to a lack of awareness, which is known to hinder protection [291]. For these reasons, we solicited ratings of device *familiarity*.

Finally, we also explored the influence of product *utility*. However, these findings were not relevant to the core of the thesis. Therefore, in the interest of brevity, *utility* results are omitted from this document.

We will now outline our four research hypotheses. In each, we compare perceptions of IoT products with those of other computers. Through this, we can establish whether smart devices differ from the baseline.

H1. Firstly, we asserted that participants would have stronger privacy concerns of IoT devices. We believed this would be due to both a ‘fear of the unknown’ [221] and the salience of media reports [131, 154, 273]. Hacks and vulnerabilities are frequently highlighted [302], and this should trigger concern. However, as previously mentioned, we expect protective behaviour to be constrained. Based on this widening disparity, the Paradox may be exacerbated. Concerns were evaluated through ratings of ‘privacy respect’. If a person deems their device to be disrespectful, they express

unease. Therefore, lower scores suggest greater concern. These ratings were on an ordinal scale from 0 (low) to 5 (high). For H1, we assert that mean *privacy scores* will be significantly less in the IoT than for the ‘non-IoT’ devices.

H2. Previous studies have criticised the usability of smart devices [143, 380]. These products often possess small screens or have few buttons. Their interfaces might also be harder to navigate, since they may require unconventional peripherals (e.g., TV remotes). Furthermore, as new mental models are necessary, people are likely to make more mistakes [247]. Therefore, we asserted that mean *usability scores* would be significantly less in the IoT than for the ‘non-IoT’ devices. This could also influence the Paradox, if individuals are less able to configure their devices [158, 230]. This might leave them prone to lax default settings [59].

H3. We expected participants would find IoT products less familiar. These devices are novel, with many being released in the last five years. They also tend to be heterogeneous, being developed by a range of vendors [345]. This should make smart products less familiar than established computers. We therefore posit that mean *familiarity scores* will be significantly less for IoT devices than the ‘non-IoT’ products. If accepted, this could have implications for the Paradox. As mental models might be misaligned, IoT products could be used in a risky manner [317]. Furthermore, users might lack the awareness to adjust their settings [247].

H4. Finally, we studied how concerns differ from actions across several devices. Concerns were gauged through the *privacy scores*, with lower metrics denoting greater unease. All ratings were made on a six-point scale, ranging from 0 to 5. Since 2/5 or 3/5 could be considered a cautious evaluation, we deemed scores under 2/5 to indicate concern. To assess actions, purchasing decisions were selected as a suitable proxy. We defined the Paradox to be expressing concern, but using the device regardless. Since we believe the IoT will be susceptible, we asserted that this disparity would be more prevalent for smart devices. Although we do not explore behaviour directly, we produce insights to inform our later studies.

We formally summarise these four hypotheses below:

1. The mean *privacy score* for IoT devices will be significantly less than the mean *privacy score* for ‘non-IoT’ products.
2. The mean *usability score* for IoT devices will be significantly less than the mean *usability score* for ‘non-IoT’ products.
3. The mean *familiarity score* for IoT devices will be significantly less than the mean *familiarity score* for ‘non-IoT’ products.

4. The proportion of individuals expressing a *disparity* for IoT devices will be significantly higher than the proportion expressing a *disparity* for ‘non-IoT’ products.

Survey design

Methodology selection. We chose an online survey for several reasons. Firstly, surveys enable the extraction of high-level discrete responses. This was appropriate at this early stage, as we wished to compare perceptions over multiple devices. If a potential disparity is identified, we can then dissect rationale through semi-structured interviews. Secondly, we first sought to analyse views at a large scale. Since interviews are more time-consuming [285], they often lead to smaller sample sizes [285]. Furthermore, to support a broad evaluation, we did not constrain our respondents to product owners. Interviews require greater effort from participants, compared to completing a quick form. Therefore, to explore perceptions at a larger scale, an online survey appeared sensible. Finally, this approach provided access to a different sample. In our first study (Chapter 4), we received opinions in specific sites. Through an online survey, we can solicit views from around the UK. If concerns are also strong in this sample, we corroborate our earlier results. Furthermore, if usability and familiarity are low, this might partially explain the Paradox. We could then continue our exploration through detailed interviews.

Refinement. The form was divided into three sections: *Demographics*, *Evaluations* and *Informed consent*. It was iteratively designed and updated over several weeks. We then undertook face validation [96] to verify whether the questionnaire was appropriate. Through this approach, the form was first reviewed by an Oxford privacy researcher. They checked whether the queries were suitable for the topic. It was then assessed by an academic experienced in questionnaire design. This person advised on bias mitigation. Based on the feedback from both parties, the instrument was further refined. The final questions are presented below in Table 5.1.

Demographics. We collected basic details to monitor for representativeness. While we also conducted demographic comparisons, these are omitted from the thesis for brevity. As in our street survey (Chapter 4), we solicited gender, age range and highest level of education. We also asked how many hours they used devices daily. We expected those acting frequently to consider products more usable and familiar.

Evaluations. In this section, participants rated the devices on: *privacy*, *usability*, *familiarity* and *utility*. In addition to providing insights, the non-privacy factors

Table 5.1: Product Perception Questionnaire

#	Demographics					
1	What is your gender?					
	<i>Male</i>	<i>Female</i>	<i>Other</i>			
2	What is your age range?					
	<i>18-25</i>	<i>26-35</i>	<i>36-45</i>	<i>46-55</i>	<i>56-65</i>	<i>66 and over</i>
3	What is your highest level of education?					
	<i>School</i>	<i>GCSE</i>	<i>A-Level/College</i>	<i>Degree</i>	<i>Masters</i>	<i>PhD</i>
4	How many hours daily do you use personal technologies?					
	<i>0-2</i>	<i>3-5</i>	<i>6-8</i>	<i>9-11</i>	<i>12-14</i>	<i>15 and over</i>
#	Evaluations (for each of six devices, all from 0/5 (low) to 5/5 (high))					
1	How usable would you rate this technology?					
2	How familiar are you with this technology?					
3	How much does this technology respect your privacy?					
4	How useful would you rate this technology?					
5	Do you own this technology?					
	<i>Yes</i>	<i>No</i>	<i>Unsure</i>			
6	Why (not)?					
	<i>(Open reply)</i>					

sought to mitigate priming [65]. To further disguise the topic, the study was advertised and presented as concerning general opinions.

Question wording is always crucial to avoid bias. We opted for ‘respecting privacy’ over ‘private’ as it had greater relevance to concerns. We believed the term ‘utility’ would be unnecessarily complex. Therefore, we chose the synonym ‘usefulness’ for the form. This should have assisted participant comprehension.

Individuals rated each factor for each category of device. As previously mentioned, we had six categories. For IoT products, we chose: *wearables*, *smart appliances* and *home automation*. The comparison group comprised: *desktops*, *laptops* and *tablets*. Again, the selection process is described later in this section. After rating devices on each factor, participants reported whether they owned the product and why. If purchases are driven by functionality, concerns might be overlooked. Furthermore, if privacy never factors into rejections, it might not be deemed important. We included an ‘Unsure’ option to avoid compelling respondents into unconfident replies. The justification enabled analysis of participant rationale. Although we considered requesting more data, this would have extended the form. Our text field sought to capture opinions while minimising the risk of response fatigue [88].

Participants evaluated each product category in turn. Since we wished to disguise the IoT/‘non-IoT’ comparison, the sequence of presentation was shuffled. Later

devices might still be rated in relation to earlier products. However, by shuffling, a single category was not disproportionately affected.

For each category, we included its name, two example devices and two images. We believed this would assist participants in identifying the technologies. However, in an attempt to minimise bias, we used a careful approach. We selected recognisable examples to encourage informed judgements. But we feared that brand predilections might influence the evaluations. For example, a Microsoft fan might criticise Apple examples, or vice-versa. Therefore, each image came from a different manufacturer.

Informed consent. To support responsible research, we received informed consent. This section outlined our data protection techniques and redressal procedures. Participants were also anonymous in the study, which should reduce the pressure from social desirability bias [138]. Furthermore, as identities were not revealed, this should mitigate non-response bias from concerned individuals [164]. Since we also received ethical approval, we believe our study was undertaken in a responsible manner.

Participant recruitment

Distribution. To complement the demographics from the street survey, we sought a diverse range of UK participants. Therefore, advertisements were placed on local and national online messaging boards. These boards included DailyInfo¹ and GumTree². Since these fora target UK consumers, we believe we surveyed members of the public.

In our recruitment message, we encouraged individuals to participate. We also included a hyperlink to the questionnaire. As previously mentioned, the topic of privacy was not advertised. We screened for adults, since these individuals should be able to give informed consent. We wished to target a broad demographic and therefore no other criteria were applied. Respondents did not have to own certain products to participate. Few individuals would possess all six devices, and this would limit our sample size. At this stage, we sought to explore high-level perceptions of IoT technologies. Since this study suggested the Paradox was likely, it was followed by in-depth interviews with product owners.

The questionnaire form was hosted on the SurveyMonkey³ portal. This site is a market-leader in online surveys, and provides secure and usable forms. Once the study was completed, the data was deleted from this service. The questionnaires were completed over a two-month period between July and August 2016.

¹<https://www.dailyinfo.co.uk/>

²<https://www.gumtree.com/>

³<https://www.surveymonkey.com/>

Compensation. To incentivise participation, we conducted a free £50 prize draw. If a respondent wished to be entered, they provided their email address at the end of the questionnaire. This was deemed preferable to the start, since we feared such disclosure might prime the topic of privacy. Since 156 individuals were in the draw, we believe their chance was fair remuneration. The compensation was optional, and participants were free to decline without penalty. Individuals also had an equal chance of winning, regardless of their product evaluations. Since we took measures to disguise the study’s purpose, we do not believe compensation biased our results.

Device selection process

Scope. In exploring our current research subquestion, we evaluated perceptions of the IoT. However, to assess its idiosyncrasies, we required a baseline group for comparison. This group comprised ‘non-IoT’ computers. Although we accept there is no direct dichotomy, we justify its usage below. The IoT is also nebulous and heterogeneous [345], ranging from self-driving cars to smart TVs. Such diversity would challenge the applicability of our findings. Therefore, we were necessitated to constrain our scope. We sought to explore the views of members of the public. As they would be unlikely to recognise industrial systems, we focused on consumer devices.

We decided to compare six product categories. This figure appeared a wise compromise between diversity and response fatigue [88]. Three were considered IoT devices and three were deemed to be less-novel. As respondents answered six questions for each category, additional products might have deterred participation.

Selection criteria. As discussed in our critique (Section 9.3), there is no strict dichotomy between the IoT and other devices. This partially due to the fact that the IoT is challenging to define [1]. A broad spectrum appears to exist from novel mobile products to traditional static computers. Furthermore, there is great heterogeneity even within device types. For example, a Fitbit wearable offers different functionality to an Apple Watch. We accept both these points. However, categories share more ‘intragroup’ similarities than between products. This is evident when comparing desktops with laptops or wearables with smart appliances. Furthermore, to determine whether the IoT presents a particular risk, we required a comparison group. Without this collection, it would be challenging to identify idiosyncrasies. We also accept that some products are portable and established, such as mobile phones. However, to categorise with greater clarity, we selected less ‘intersectional’ technologies.

We specified three criteria to select our devices: *novelty*, *ubiquity* and *autonomy*. These criteria were chosen as we deemed them indicative of IoT products. Therefore,

their inverse could be used to designate alternative technologies. We will now justify the selection of these criteria.

In terms of *novelty*, the IoT has only matured in the past decade [311]. Indeed, many archetypal products, such as smartwatches, gained popularity in 2012⁴. In contrast, traditional computers are more established. Laptops emerged over 30 years ago⁵, while early desktops date to the 1960s⁶. With the latest trend being for Internet-connected gadgets, novelty appeared a sensible differentiator.

Ubiquity is considered a key trait of many IoT products [249]. Smart devices can attach to our bodies [364] or pervade our personal environments [264]. This allows them to offer functionality not supported by traditional computers. In contrast, desktops are static and often reserved for clerical tasks. Although laptops can move from room to room, they lack the pervasiveness of home automation systems. The IoT seeks to blur the relationship between the physical and virtual worlds [146]. Therefore, ubiquity was an apt criterion.

Finally, *autonomy* allows products to interact with their surroundings [216]. This might be through sensor monitoring, as found on smartwatches and home automation systems. Alternatively, actions may be triggered remotely, as is common with smart appliances [231]. In contrast, computers tend not to possess physical sensors. This can limit their capability at environmental monitoring. As autonomy is a common IoT characteristic [308], it was selected as our final criterion.

Device selection: IoT. Through these three criteria, we sought to identify IoT technologies. To achieve this, candidate products were plotted against the axes. We deemed quantitative measures to be impractical, particularly when evaluating *ubiquity* and *autonomy*. Therefore, the plotting was undertaken loosely. An illustration can be found in Figure 5.1. As shown in the top right, IoT devices tend to be novel, ubiquitous and autonomous. The bottom left reflects how other products are less represented by these criteria. We continue by discussing our chosen devices.

Wearables, considered to be smartwatches or fitness trackers, were placed in the IoT grouping. These devices met the novel criterion since they only achieved maturity in the last five years⁷. By their nature they are also ubiquitous, being worn on an owner's wrist. Finally, wearables often make use of autonomous sensors. They might track running distance or heart rate. As such products are commonly included in IoT definitions [265, 329], we are confident in their categorisation.

⁴<https://www.wareable.com/smartwatches/smartwatch-timeline-history-watches>

⁵<http://www.laptop-lcd-screen.co.uk/shop/historyofthelaptopcomputer.asp>

⁶<https://www.computerhope.com/issues/ch000984.htm#desktop>

⁷<https://www.wareable.com/smartwatches/smartwatch-timeline-history-watches>

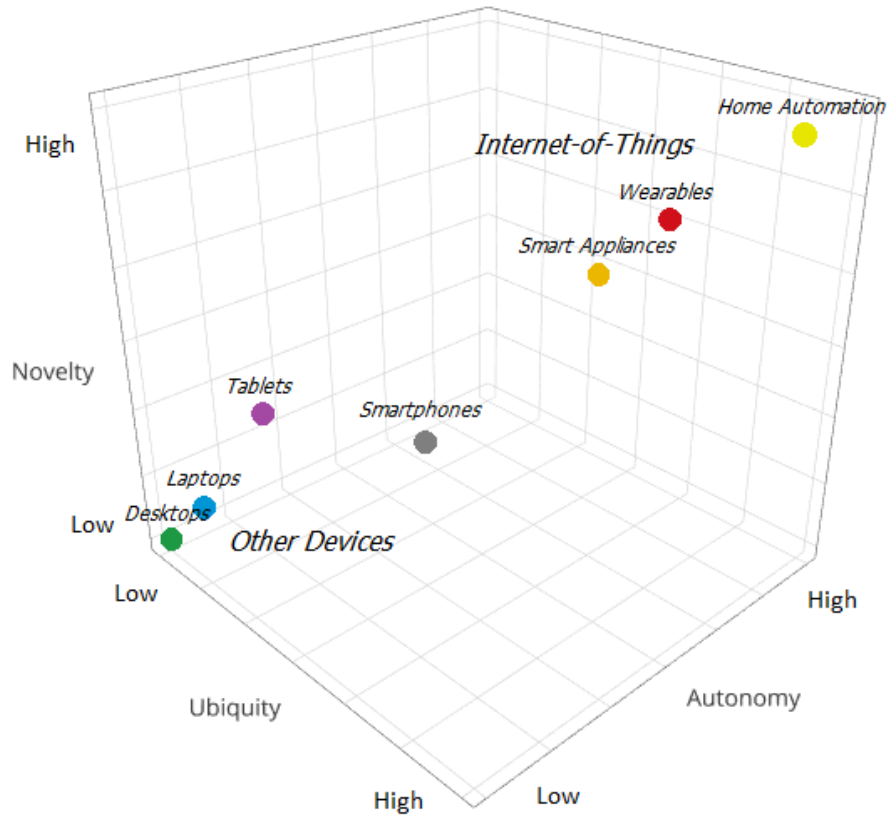


Figure 5.1: Device Category Selection

Smart appliances, such as smart TVs and connected fridges, are also novel devices. These technologies were first released ten years ago [84], but have only achieved recent popularity. Although many lack mobility, they express ubiquity. Appliances pervade personal environments and perform non-clerical tasks. These systems can also act autonomously, reacting to environmental factors [63]. Since smart appliances are often considered archetypal of the IoT [184], their selection is apt.

Home automation systems, such as Hive and Google Nest, are still relatively novel [368]. These platforms have seen increased adoption over the last five years [234], but are still rare. Although the systems are static, they are ubiquitous throughout the home. Interfaces might be found in multiple rooms, and the system can affect the whole house. Finally, by definition, home automation systems offer great autonomy. Based on these factors, we deemed them an IoT product.

Device selection: Non-IoT. While we oppose a strict dichotomy between the IoT and computers, ‘non-IoT’ was chosen as an accessible abbreviation. To assess

the idiosyncrasies of smart devices, we sought products which differed. Therefore, to select these technologies, we considered devices near the opposite side of our axes.

Desktops are far from novel, having been popular for several decades. The platform originated in the 1960s⁸, with such machines widespread since 2000⁹. These products also lack ubiquity, with them often placed in home offices. They are static by nature and are frequently inconvenient to move. In addition, desktops lack the autonomy of the IoT. Although they can be accessed remotely, local peripherals are commonly used. Finally, they rarely possess physical sensors, limiting interaction with their surroundings. Based on these criteria, desktops are not an IoT product.

Laptops were developed more recently, but were still released over 30 years ago¹⁰. While they are more portable than desktops, they are not considered an IoT device [123, 325]. Furthermore, they rarely interact with their surrounding environment. Finally, laptops lack autonomy as they tend not to possess physical sensors. Therefore, these devices belong in the ‘non-IoT’ grouping.

Tablets are closer to the intersection, but differ from IoT products. While smart devices have gained recent popularity, tablets date from the 1960s¹¹. The Amazon Kindle was a breakthrough product and this was launched over a decade ago¹². In terms of ubiquity, the devices are certainly portable. But they tend not to interact with their physical environment. Finally, tablets rarely possess the sensors of other IoT devices. They also lack autonomy and are usually driven by local instruction. As they are not conventionally classified as an IoT product [123, 325], we are confident in our categorisation.

Chosen categories. Based on this selection process, we selected six categories for comparison. IoT devices consisted of: *wearables*, *smart appliances* and *home automation*. We also considered ‘drones’, also known as Unmanned Aerial Vehicles (UAVs). However, we deemed these products to be less representative than the above systems. The ‘non-IoT’ products were: *desktops*, *laptops* and *tablets*. The latter three technologies have also been judged distinct from the IoT [123, 325]. We therefore have confidence in this division.

⁸<https://www.computerhope.com/issues/ch000984.htm>

⁹www.statista.com/statistics/289191/household-penetration-of-home-computers-in-the-uk/

¹⁰<http://www.laptop-lcd-screen.co.uk/shop/historyofthelaptopcomputer.asp>

¹¹<https://www.techradar.com/news/mobile-computing/10-memorable-milestones-in-tablet-history-924916>

¹²<https://www.pocket-lint.com/gadgets/news/137303-amazon-kindle-history-kindle-to-the-kindle-oasis>

5.3 Results and Discussion

Participants

We now consider the results from our online survey. We received 170 responses to the questionnaire. Of these participants, 57.1% were male and 42.9% were female. This is in contrast to the street survey (57.1% female, 42.9% male), providing an exploration of a different UK demographic. In terms of age, 50% were 26-35 and 25.9% were between 18 and 25. This was reflected by our estimated mean age of 32.0, indicating a relatively-young sample.

Our respondents appear highly-educated: 35.9% had a Master's degree, while another 38.8% reported a Bachelor's qualification. In terms of daily usage, the estimated mean was 7.4 hours. This implies that our participants should be accustomed to digital devices. We expect this will contribute to informed evaluations.

Qualitative analysis

In addition to studying participants' evaluations, we also explored whether they owned each device. If a person expresses concern yet purchases a product, their claims and behaviour may not align. Moreover, by analysing ownership, we can ascertain which technologies are most appreciated.

When studying ownership, we also request a qualitative justification. This brief description serves two primary purposes. Firstly, we can explore the decision-making rationale of our respondents. The factors which influence their purchases are likely to also affect their behaviour. For example, if users are driven by functionality, protection might be deemed less important. Secondly, we can investigate whether privacy is salient in these decisions. If a person expresses strong concern yet fails to consider privacy, their behaviour might not be aligned.

Coding frame creation. To analyse these qualitative justifications, we followed a robust approach. We adopted the inductive analysis [323] detailed in Section 3.3. To begin this process, we read our responses several times. Once we were familiarised with the data, we continued by annotating draft labels. Through this approach, general themes began to emerge. For example, some purchased desktops to support their programming activities. In contrast, many rejected wearables due to their excessive cost. These topics were then used to form six coding indices. One index was produced for each device category, since justifications varied based on the product. Themes were divided into subthemes as we developed rich hierarchies. The refined

indices then became our six coding frames. We added strict definitions, seeking to reduce ambiguity. If additional topics emerged during coding, these tables were further updated. Our final frames can be found in Appendix A.1.

Coding approach. Participants often gave multiple justifications within a single response. We did not wish to select a single reason, since this would reduce the depth of our analysis. To support a rich exploration, each response could be assigned multiple themes. The themes were still independent as they all concerned distinct topics. Often, this led to more codes being assigned than the number of respondents. Therefore, theme proportions were reported in the ‘comment’ fashion (Section 3.3). On the few occasions that we discuss participant totals, it is clearly signposted.

Privacy concerns

Omnibus. We began by investigating privacy concerns, since these relate to one half of the Paradox. If a product receives low evaluations, it indicates that respondents express concern. To explore this factor, we compared the mean *privacy scores* of each device. To analyse whether concerns vary based on product, we first undertook a Friedman test. *Privacy scores* indeed differed significantly ($X^2(5) = 70.470$, $p < 0.001$), suggesting some devices provoked greater unease.

To identify which devices were responsible, we then conducted Wilcoxon Signed-Rank comparisons. We made 15 pairwise analyses, since each product was compared against its five alternatives. Bonferroni correction was then applied to mitigate the risks from the Multiple Comparison Problem [259]. This resulted in an adjusted significance threshold of $\alpha = 0.003$ ($0.05 \div 15$). We now discuss how products were considered in terms of privacy concerns.

Comparisons. Laptops were regarded as the most privacy-respecting device, receiving a mean rating of 3.27/5. This implies that participants were least concerned about these computers. The ratings might owe to the fact that many privacy-protective tools are available. The *privacy score* was significantly greater ($Z = -4.043$, $p = 0.001$, $d = 0.652$) than that of desktops, which scored 2.99/5. Tablets were rated third-most privacy-respecting ($\bar{x} = 2.76/5$), though the difference from desktops was not significant. The three ‘non-IoT’ products appeared to provoke the least concern. This suggests that smart devices might increase this Paradox component.

Home automation came fourth, receiving a mean *privacy score* of 2.64/5. While this differed from laptops ($Z = -4.641$, $p < 0.001$, $d = 0.762$), it did not from desktops or tablet devices. Smart appliances provoked the second-largest concern ($\bar{x} = 2.58/5$), but also did not differ from tablets. These products might be worrying due to the

media reports of eavesdropping [50]. Finally, wearables received the lowest scores, rated only 2.31/5. Although not worse than appliances, it differed from laptops ($Z = -6.580$, $p < 0.001$, $d = 1.169$), desktops ($Z = -4.767$, $p < 0.001$, $d = 0.786$) and tablets ($Z = -3.872$, $p < 0.001$, $d = 0.622$). Wearables can be used to track their owner’s movements, and this prospect might trigger unease [125]. Based on the sequence of ratings, the IoT appears to provoke the greatest concern.

H1. In our first hypothesis, we asserted that mean *privacy scores* would be lower for the IoT than for ‘non-IoT’ products. We found that the smart device average was only 2.51/5, compared to 3.01/5 for ‘non-IoT’. Through significance testing, our hypothesis was accepted ($Z = -5.151$, $p < 0.001$, $d = 0.86$). This implies that IoT products, particularly wearables, provoke the greatest privacy concern. If protective behaviour is not strong, the Paradox might be prevalent.

Usability evaluations

Omnibus. We moved on to consider the mean *usability scores*. If a product is challenging to operate, this might impede privacy-protective behaviour [388]. To compare ratings at an omnibus level, we conducted another Friedman Test. The scores differed significantly based on device ($X^2(5) = 292.832$, $p < 0.001$), implying that some products were thought less usable. In order to determine the sequence, pairwise comparisons were then undertaken. Bonferroni correction was applied as before, resulting in a significance threshold of $\alpha = 0.003$.

Comparisons. Laptops were considered the most usable product ($\bar{x} = 4.55/5$), in addition to being the most privacy-respecting. They were rated significantly greater than desktops ($Z = -5.869$, $p < 0.001$, $d = 1.008$), who were again in second place ($\bar{x} = 3.96/5$). These high scores might be due to the large screens and sizeable keyboards. Tablets were next in third position, receiving a mean rating of 3.81/5. Based on this sequence, IoT products received the lowest evaluations. If they lack usability, individuals might be impeded in their privacy protection [158, 230].

IoT devices all received lower evaluations. Smart appliances came fourth ($\bar{x} = 3.12/5$), with ratings differing from those of tablets ($Z = -5.550$, $p < 0.001$, $d = 0.941$). Although some of these machines have large screens, they may lack standard input peripherals (e.g., TV remotes). Home automation had an average of 3.00/5, but this did not differ from connected appliances. Wearables again received the lowest scores, significantly less than the above products (vs Smart Appliance: $Z = -4.678$, $p < 0.001$, $d = 0.769$; vs Home Automation: $Z = -3.577$, $p < 0.001$, $d = 0.571$). Smartwatches possess small screens and therefore might be challenging to use

[52]. This implies that not only do they provoke the greatest concern, they also lack usability. If this constrains protective behaviour, the Paradox might be prevalent.

H2. We now explore our second hypothesis, which posited that IoT products would receive lower mean *usability scores*. While smart devices were rated only 2.88/5, the other group had an average of 4.11/5. Through a Wilcoxon Signed-Rank test, we found IoT products were indeed considered less usable ($Z = -10.332$, $p < 0.001$, $d = 2.598$). This ‘huge’ effect size [333] emphasises the divide between the technologies. If the IoT is challenging to operate, protective behaviour might be impeded [158, 230].

Familiarity evaluations

Omnibus. We have already found that IoT devices both provoke the greatest concern and possess the least usability. If these products are also considered unfamiliar, this could further constrain protection [317]. As before, we used a Friedman test to ascertain whether ratings differed significantly. This was found to be true ($X^2(5) = 512.668$, $p < 0.001$), so we proceeded to pairwise comparisons. We applied Bonferroni correction ($\alpha = 0.003$) to again mitigate the Multiple Comparisons Problem.

Comparisons. Laptops received the highest ratings ($\bar{x} = 4.72/5$), suggesting they were familiar to participants. Since they were released several decades ago, this was far from surprising. They were considered significantly more familiar than desktops ($Z = -3.367$, $p = 0.001$, $d = 0.535$), which were rated 4.60/5. Although desktops were launched even earlier, they might have lost popularity due to their size and immobility. They were rated higher than tablets ($Z = -6.514$, $p < 0.001$, $d = 1.153$), which again occupied third position ($\bar{x} = 3.95/5$). As before, IoT products all received the lowest evaluations. Since they have been released more recently, this likely affected their familiarity.

Smart appliances were rated 2.64/5, likely because they have not yet achieved widespread adoption. This rating was significantly less than the tablet evaluation ($Z = -8.222$, $p < 0.001$, $d = 1.625$). However, scores were greater than those for wearables ($Z = -4.471$, $p = 0.001$, $d = 0.73$), which again were low (1.95/5). With smartwatches lacking usability and familiarity, protective behaviour might be impeded [158, 230, 317]. Home automation systems were considered least familiar (1.45/5), with a lower mean than wearables ($Z = -3.377$, $p = 0.001$, $d = 0.536$). Since these systems are expensive and challenging to install, this might limit their popularity.

H3. Our third hypothesis asserted that IoT products would be less familiar than other computers. The *familiarity scores* appeared to support this, with the

smart device mean (2.01/5) being much lower (4.42/5). Following a Wilcoxon Signed-Rank test, we accepted the hypothesis ($Z = -11.103$, $p < 0.001$, $d = 3.249$). This ‘huge’ effect size [333] demonstrates the difference between the technologies. With large effects found for each hypothesis, it suggests that IoT products are judged to be different. If concern is strong but protection is impeded, the Paradox might be prevalent. This is particularly the case for wearables, which received low scores across the three aforementioned factors.

To illustrate the differences in ratings of IoT and other technologies, we produced the below figure (Figure 5.2). As displayed, the former group provoked the greatest degree of privacy concern. This might increase the first component of the Paradox. They were also considered less usable and (less) familiar, which might constrain protective behaviour [317]. This could decrease the Paradox’s second component. Therefore, as the IoT proliferates, this disparity might expand.




Figure 5.2: Device Mean Factor Ratings

Ownership decisions

We moved on to consider purchasing decisions. As IoT devices were both novel and less familiar, we expected them to be possessed by fewer participants. This percentage was denoted as the *ownership frequency*. However, it was the justification to which

we paid particular attention. Respondents supplied a one-line rationale, which was analysed through our qualitative techniques. Participants might purchase products for many reasons, including style, usability or functionality. They might also reject due to several factors, including expense and unfamiliarity. While our respondents appear concerned about privacy, this might not encourage cautious behaviour. If a product is purchased despite these claims, the two factors might not align.

Non-IoT ownership. We now discuss *ownership frequencies* and participant justifications. ‘Comment’ proportions are reported since many reasons might be given for a single response. As suggested by the familiarity ratings, *laptops* were the most popular product. They were owned by 92.9% of participants, reflecting the prevalence of this device. The most frequent justification was its use for work (22.4%). Portability was also appreciated (22.0%), as was the entertainment that it supports (8.3%). This was not surprising, as laptops are useful devices. The most frequent rejection rationale was that users preferred other products (58.3%). Of the 7 who responded in this manner, 4 opted for desktops and 3 preferred tablets. Interestingly, not a single privacy concern was expressed. It would appear that this factor rarely influences laptop purchases.

Tablets were next-most popular, owned by 65.9% of our participants. Purchases were again justified by portability (21.0%), likely due to the compact size. Individuals also appreciated the quick convenience (8.9%) and the usability (7.0%). Rejection was primarily blamed on other device use (49.5%), with 46.2% of these favouring smartphones and 46.2% preferring laptops. While they are too large to fit in a pocket, they often lack usable keyboards [85].

Only a single individual criticised privacy, suggesting these views were not widespread. They disliked that users lacked complete control of the product’s security. Interestingly, the user still chose to purchase the tablet. This implies that concerns were not severe enough to encourage rejection. Representative quotes, along with participant number, are displayed below.

“Convenient to use handheld devices, but not 100% give user full control and customisation of privacy.” (#30).

Desktops were next in popularity, possessed by 52.4% of respondents. Ownership rationale was varied, with 34.3% praising desktop applications. Within this, 22.9% enjoyed entertainment features and 16.7% liked gaming. Of the 104 rejection comments, the most-common was a preference for other devices (45.2%). Laptops were chosen in 85.1% of these cases, likely since they are more portable. We then considered

whether privacy was a salient factor. As with laptops, not a single person criticised this matter. Across our ‘non-IoT’ products, privacy appears a rare consideration.

IoT ownership. We now consider the rationale behind IoT purchases. *Smart appliances* were fourth-most popular, owned by 41.8% of participants. This likely owes to the prevalence of smart TVs, with 25.7% of justifications mentioning this functionality. Participants also appreciated the entertainment value (12.9%) and efficient convenience (11.9%). In terms of rejection, individuals were discouraged by high prices (24.5%) and a lack of necessity (37.3%).

Participants appeared somewhat aware of privacy issues. Five individuals rejected based on this factor, with the main issue being perceived surveillance (4.5%). This was likely due to the media reports concerning eavesdropping televisions [50]. While it is encouraging that privacy is considered, price was mentioned 5x more often. As markets mature, technology prices tend to naturally fall [334]. This suggests that appliances might increase in popularity.

“They may even know or predict what my lifestyle is. It is like a big brother always spying on you.” (#30).

Wearables were far less popular, only possessed by 20.6% of participants. Of the reasons justifying ownership, 44.2% praised the exercise-monitoring functionality. This is understandable since fitness trackers are a common wearable. Reasons for rejection were highly varied. The most popular rationale (53.1%) was that the technology was unnecessary. Other justifications included the cost (17.9%) and a smartphone preference (9.3%).

As previously mentioned, wearables provoked the greatest privacy concern. However, only 3.1% of comments considered this factor. These individuals disliked the notion of being continuously monitored. However, a large number of participants appeared to appreciate this exercise tracking. This demonstrates tension between privacy and functionality, potentially contributing to the Paradox. Although privacy was mentioned, price was cited over 5x more frequently (17.9%). Therefore, as the market matures, these devices might increase in prevalence [334].

“Because the wearer is never unreachable.” (#72).

Home automation systems were least-common, owned by only 12.4% of participants. Of the few who purchased these technologies, 16.0% justified this on useful functionality. Another 8.0% considered them usable, while 8.0% felt they saved

money. In terms of rejection, the most popular reason was price (22.4%). Another 22.4% found them unnecessary, implying they have not gained broad appeal.

When considering *privacy scores*, participants expressed concern. However only 1.1% of comments referred to the topic. They cited ‘privacy’ in general, rather than outlining specific issues. In contrast, cost was blamed over 20x more often. As before, while concern is claimed, it rarely appears to influence decisions. If true, the IoT Paradox might be prevalent.

“I would have privacy/security concerns and do a lot of homework before installing anything, if I wanted it.” (#116).

The Privacy Paradox

Our survey does not explore the use of protective settings. Therefore, we reserve a detailed Paradox analysis to the next study. However, we can receive an initial insight through considering product purchases. If individuals claim concerns but buys it regardless, a disparity might exist. Our final hypothesis (H4) asserted that such disparities would be significantly more frequent for IoT devices. To explore this, we begin by considering each product individually. Since we analyse views on a participant level, we report ‘participant’ proportions.

Non-IoT. In these cases, any disparity between concern and behaviour was rare. For laptops, only 13/158 (8.2%) bought the product while expressing concern. This is likely because their mean *privacy score* was the highest, and hence few were worried. The percentage was slightly higher for tablets, with 12/112 (10.7%) acting in this manner. These products offer convenient functionality, and this might overcome concerns. Desktop percentages were also quite low: only 7/89 (7.9%) displayed a disparity. We took a mean aggregate average, finding the ‘non-IoT’ rate was 9.0%. While this does not consider protective behaviour, the issue does not appear common.

IoT. We now move on to assess smart devices. For home automation systems, 2/21 (9.5%) displayed a disparity between concern and behaviour. While these products trigger some worries, users might care more about convenience. Smart appliances performed even worse, with 9/71 (12.7%) acting in this manner. Such devices appeared popular with the public, being owned by over 40% of our respondents. However, since they are suspected of eavesdropping [50], they can trigger privacy concerns. Finally, wearables contributed to this disparity most frequently (6/35, 17.1%). Since they also lacked usability and familiarity, protective behaviour might be impeded [158,230,317]. When we calculated an IoT mean, 13.4% displayed this issue.

H4. For our final hypothesis, we asserted that this disparity would be more common in the IoT. Although the proportion appeared greater than the ‘non-IoT’ rate, the difference was not significant ($p = 0.150$). Although the percentage was higher for IoT devices, there were comparatively few owners. For example, while 158 respondents had laptops, only 21 possessed home automation systems. In larger samples, it might be possible to find significance. As previously noted, purchases were not a perfect proxy for behaviour. When analysing protective actions, usability and familiarity might have an influence [158, 230, 317].

Through H1, we found that the IoT provokes the most concern. H2 and H3 showed that these devices are considered less usable and less familiar. Moreover, based on behavioural rationale, privacy appears rarely taken into account. Therefore, we believe that the IoT Paradox will be prevalent, particularly in wearable environments.

5.4 Implications

Concerns. To address our second research subquestion¹³, we compared perceptions towards a variety of devices. This was undertaken through an online survey with 170 participants. We found that IoT products provoked the greatest concern. Wearables were deemed the worst, receiving significantly lower scores than ‘non-IoT’ technologies. This suggests that the concern component of the Paradox was high.

Constraints. Participants considered IoT products to be significantly less usable. If a device is complex to operate, privacy features might be neglected [247]. These products were also rated significantly less familiar, with this posing additional constraints [317]. Such issues might impact the behaviour component of the Paradox. If these factors diverge, the issue might be prevalent. As highlighted, this study used purchasing decisions as a proxy. To assess protective behaviour in depth, we conduct semi-structured interviews in the next chapter.

Wearables. Wearables may be susceptible to the Paradox. They provoked the greatest privacy concerns, while also being considered least usable. Furthermore, they were rated as lacking in familiarity. Since concerns are strong and behaviour might be impeded [317], wearables could pose a risk. This supports claims that mobile computing can threaten rational privacy decisions [47]. When studying justifications, cost appeared more influential than privacy. If smartwatches mature and their prices decrease [334], these devices could proliferate.

¹³SQ2: *How do perceptions differ between IoT and other devices?*

Rationale. Although strong concerns were expressed, privacy was rarely factored into purchasing decisions. Participants appeared far more driven by price, features and usability. Indeed, elements of functionality were the largest factor for all IoT devices. If individuals become preoccupied, they might overlook privacy [103]. Furthermore, when the topic lacks salience, users might not protect their data [188]. To mitigate a Paradox, perhaps the principle should be highlighted.

Next steps. Thus far, we have addressed our first two research subquestions. In Chapter 4, we confirmed the existence of the Paradox in the UK. Our online survey now enabled a comparison of perceptions across several devices. Based on our findings, the disparity might be common on IoT devices. Through purchasing decisions, we conducted an initial exploration of behaviour. However, to study the matter at greater depth, we require rich data. This can be provided through informed discussions of owned devices. Concerns can be deconstructed based on reactions to contextualised questions. The use of protection can be assessed, alongside rationale for this behaviour. Therefore, to explore the prevalence of the Paradox, we conduct semi-structured interviews.

Chapter 6

How and why does the Privacy Paradox differ between IoT and other devices?

6.1 Introduction

Background

In our first study (Chapter 4), we confirmed that the Privacy Paradox existed in the UK public. Through our second work (Chapter 5), we found that IoT privacy concerns appear to be greater than those towards ‘non-IoT’ devices. This implies that the concern component of the Paradox might increase. We also discovered that IoT products were considered significantly less familiar and (less) usable. Previous work has suggested that such issues can constrain protection [158, 317]. Therefore, these factors might reduce the behaviour component. This divergence might lead to the Paradox increasing in prevalence.

While our studies have analysed the Paradox at a high-level, we have not explored rich rationale. In the previous chapter (Chapter 5), purchasing was used as a proxy for behaviour. This was appropriate in our high-level survey of opinion. However, to address our third research subquestion (*How does the prevalence of the Privacy Paradox differ between IoT and other devices?*), protective behaviour must be considered.

The focus of the thesis is to reduce the disparity, in addition to studying it. To design mitigative approaches, we must understand why individuals act in this manner. Therefore, we continue by conducting in-depth interviews with product owners. This supports us in addressing our fourth research subquestion: *Which factors contribute to the Privacy Paradox in the IoT?*

Motivation

Previous work has criticised studies for comparing abstract concerns with practical behaviour [369]. Since privacy is contextual [278], users might value the concept in principle but take little action. To explore the Paradox with nuance, discussions will be contextualised around defined devices.

Although the online survey extracted opinions, only 3/170 participants (1.8%) owned all the products. This was sufficient in this case, as we compared general concerns against purchasing decisions. However, for an detailed dissection of behaviour, we now recruit device owners. This should support the extraction of informed views.

To explore the influence of the IoT, we require a baseline group for comparison. Therefore, to establish consistency, we use the same six categories from the previous chapter. Again, while we accept there is no IoT-computer dichotomy, the products were selected through the robust process. Through recruiting owners and comparing their responses, we can ascertain whether the Paradox prevalence differs.

6.2 Methodology

Hypotheses

Metrics. We begin by outlining the metrics we receive from this study. Our interviews seek to explore both privacy concern and protective behaviour. As will be described in the Question Design subsection, privacy concerns were gauged through contextualised questions. This approach is similar to that of ‘scenarios’, as deemed successful in previous studies [2, 133, 232]. Our selection criteria are outlined in Section 3.5. Rather than querying concern directly, responses were requested to four incidents. Since they did not mention ‘privacy’, this should avoid priming the topic. This is important, as social desirability bias [138] can inflate concerns.

To assess the matter fairly, it was crucial that these scenarios had feasibility. This was achieved by considering a detailed threat model, as described in Section 3.4. Since we compare a wide range of devices, we opted for the *general* construct. Based on the responses to four questions, to be introduced in Interview Question Design, we formed the *concern score*. This metric ranged from 1/5 (low) to 5/5 (high). Score generation is outlined near the beginning of Section 6.3.

To explore the Paradox directly, we must understand protective behaviour. Our analysis was undertaken by considering the frequency of three actions. The criteria for their selection is outlined in Section 3.5. These techniques, also introduced in

Interview Question Design, all improve the privacy posture of an individual. If a person wished to protect their device, they could take these actions. We formed the *behaviour score* from responses, with the process described in Section 6.3. This metric ranged 1/5 (low) to 5/5 (high).

H1. We now turn to our hypotheses. In Chapter 5, we found IoT devices were considered less usable and (less) familiar. Based on our aforementioned rationale, we believe protective behaviour might be constrained. Therefore, we assert that the mean IoT *behaviour score* will be significantly less than that for ‘non-IoT’ products. If accepted, this might imply that smart devices constrain protection. Wearables are expected to perform poorly, since they received low ratings in the previous study.

H2. An absence of protection does not alone indicate the Paradox. If concerns are also low, then opinions and behaviour are commensurate. However, we do not expect IoT users to lack concern. Indeed, Chapter 5 demonstrated that these products can provoke unease. This might be through their functionality or a ‘fear of the unknown’ [221]. Even if data is considered innocuous, invasions might be rejected on principle. As we posit that IoT behaviour will be less-protective, the disparity should increase. With *concern scores* and *behaviour scores* rated out of five, a *disparity* was defined as a two-point gap. While a one-point gap would not signify a dissonance, we believed three points to be excessive. For our second hypothesis, we assert that this disparity will be significantly more frequent for the IoT than ‘non-IoT’ products. If Paradox prevalence does differ, this might be due to the devices themselves. By exploring the matter across a range of products, we address our third research subquestion.

We formally summarise our two hypotheses below:

1. The mean *behaviour score* for IoT devices will be significantly less than the mean *behaviour score* for ‘non-IoT’ products.
2. The number of individuals expressing the *disparity* for IoT devices will be significantly higher than the number expressing the *disparity* for ‘non-IoT’ products.

While not a falsifiable hypothesis, we also seek to determine those factors that contribute to the Paradox. This is achieved through a robust analysis of discussions. If the disparity is indeed prevalent in the IoT, we can begin addressing the issue. Through analysing which justifications are most frequent, we can design countermeasures. In this manner, the findings from this study will directly support subsequent chapters. By identifying contributory factors, we address our fourth subquestion.

Study design

We solicited concerns and behaviour through semi-structured interviews. These discussions were grounded in a defined context, since all questions related to a particular owned device. In this manner, we adopted the important *principle of compatibility* [12]. This supports fair comparisons between concerns and behaviour [369].

Participant groups. To explore how the Paradox relates to the IoT, we required a baseline group for comparison. Therefore, we recruited two sets of participants. The former possessed IoT devices, while the latter owned less-novel products. This second group acted as a control to enable smart device analysis. Without this sample for grounding, it would be harder to determine the IoT’s idiosyncrasies. Furthermore, since we expect the Paradox to be exacerbated, we required a disparity baseline.

For consistency, we used the category definitions from the previous chapter (Chapter 5). Therefore, the IoT group had owners of: *wearables*, *smart appliances* and *home automation systems*. ‘Non-IoT’ participants owned: *desktops*, *laptops* and *tablets*. We again accept that technologies are not in a strict dichotomy, with this discussed in our critique (Section 9.3). To contextualise our analyses of concern and behaviour, all questions focused on a single device. When individuals possessed multiple products, they were asked to select one they used frequently. Through this, we sought to explore informed responses.

Interview scheduling. The 20-minute discussions took place over a three-week period. Interviewees selected a timeslot themselves, with sessions available between 9am and 6pm, Monday to Friday. We accept that daytime scheduling might have limited the participation of certain demographic groups. However, it was preferable to evenings or weekends, where individuals might have been reluctant to donate their time. Furthermore, with interviews conducted in an unfamiliar location, participants might feel uneasy after dark.

Interview procedure. Our discussions contained only two individuals: a researcher and a participant. To control for external factors, the same researcher conducted all the interviews. All respondents faced the same questions, with only the device name customised in our between-subjects format.

The interviews were conducted in a meeting room within the university. This was considered preferable to Skype or phone discussions, which can challenge rapport [330]. Interviewees first read an information sheet, detailing data storage and redressal procedures. They then completed an informed consent form, explicitly agreeing to audio recording. By recording the discussions, we sought to encapsulate the richness of data. As no participants took exception to this point, we believe it was appropriate.

Participants then completed a brief demographics sheet, enabling the monitoring of representativeness. As before, we omit discussion of comparisons for the sake of brevity. Establishing consistency with the previous studies, we solicited gender, age range and highest education level. We also requested a self-evaluation of how technologically-literate each person believed themselves to be. If we receive a literate sample, this might place a minimum bound on Paradox prevalence [48, 114].

Once the discussion was completed, respondents were given their compensation (outlined below). Since the privacy theme had been disguised, we then debriefed them on the interview's topic. Finally, they were guided from the meeting room to the exterior of the building.

Participant recruitment

Distribution. Recruitment was undertaken for two weeks and then the interviews proceeded for five (January and February 2017). To explore the prevalence of the Paradox, we recruited a distinct sample. This complemented the multi-site sample from the street surveys and the remote group from the online questionnaires. Since our interviews were conducted physically in the university, we recruited from the local area. However, to avoid student samples, we advertised outside the institution. Recruitment was undertaken through a local messaging board, DailyInfo. This message board was non-academic and popular with the general public. Therefore, our findings are not limited to student perceptions. Although such demographics are necessary in some studies (Chapter 8), we sought the rationale of a varied sample.

Screening. Applicants were asked to email the organisers if they met two conditions: that they were willing to participate in interviews, and that they owned a device from one of our six categories. While this ownership was not verified, participants discussed their product for 20 minutes. It is unlikely that such an in-depth conversation would be wholly fabricated. The only other criteria were that the individual was an adult and able to give informed consent.

Response bias mitigation. It was crucial that the most privacy-conscious individuals did not avoid our study. To reduce this non-response bias, we highlighted that data would be stored securely and anonymously. Since priming can inflate concerns [65], we also sought to disguise the topic of privacy. Furthermore, if the study purpose is known, demand characteristics can introduce bias [287]. Therefore, the interviews were advertised as concerning 'modern technology' and privacy was not mentioned. The question sequence was also adjusted to mitigate order effects. Due to these approaches, we believe biases were reduced.

Compensation. Participants were compensated for their time with £10 in cash. Since they were required to visit the university and devote 20 minutes of their day, we felt this was fair remuneration. This is particularly the case since the meeting room was outside the city centre. Interviewees were under no obligation to accept the funds, and could reject it without penalty. Although compensation can disproportionately encourage those with lower incomes [160], £10 was not a large sum. Since the topic and aims of the study were disguised, we do not believe this affected our findings.

Interview question design

Structure. Our interviews followed a semi-structured approach, allowing for topics of interest to be further explored [46]. This was preferred to structured format, where we would be less able to investigate user rationale [285]. We also avoided an unstructured approach, where fair comparisons might be challenging [285].

Refinement. In seeking to enhance clarity and mitigate biases, the questions were refined over several weeks. As in the previous chapter, we then received face validation [96] from two individuals. A privacy-focused colleague verified that the topic was addressed. A questionnaire expert, also from our institution, identified potential biases. Once the design was updated, we conducted a small pilot study within our research team. We found our original question sequence led to order effects. After moving our behaviour queries later, privacy appeared better-disguised.

Design. Interview questions were of four main types: *General*, *Concern*, *Behaviour* and *Privacy Paradox*. The queries can be found below in Table 6.1. They were designed to be broad, soliciting open-ended rationale from the general public. While a prescriptive approach may have channelled responses, we would have constrained the diversity of replies. Since we sought rich data to dissect rationale, this would have been counter-productive.

General questions. If the topic of privacy was immediately discussed, participants might become primed and adjust their responses [65]. Therefore, these questions both disguised the matter and solicited general opinions. If a participant mentions privacy at this unprompted stage, then the concept might be valued.

For the sake of brevity, we only highlight the questions relevant to our discussion. Near the start of this section, we solicited a general opinion. We also asked participants what they liked most and least about their product. If privacy arose in these unprompted questions, it suggests an interest in the topic. We next queried why they thought others might not purchase their device. Even if a person is not concerned,

Table 6.1: Semi-Structured Interview Questions

#	Demographics
1	With which gender do you most identify? <i>Male</i> <i>Female</i> <i>Other</i>
2	What is your age? <i>18-25</i> <i>26-35</i> <i>36-45</i> <i>46-55</i> <i>56-65</i> <i>66 and over</i>
3	What is your highest level of education? <i>School/GCSE</i> <i>A-Level/College</i> <i>Degree</i> <i>Masters</i> <i>PhD</i>
4	How technologically-literate would you rate yourself? <i>1/5</i> <i>2/5</i> <i>3/5</i> <i>4/5</i> <i>5/5</i>
#	General Questions
1	What device did you buy?
2	How much did you know about your X before you bought it? Why?
3	What is your opinion of your X? Why?
4	What do you like most about your X? Why?
5	What do you like least about your X? Why?
6	Have you considered upgrading your X to a newer version? Why?
7	Why do you think some people don't buy Xs?
8	How often do you completely turn off your X? Why?
9	Who do you think has access to your X's data?
#	Concern Questions
1	How would you feel if someone deleted your X's data without your permission? Why?
2	How would you feel if someone shared your X's data without your permission? Why?
3	How would you feel if someone monitored everything you do on your X? Why?
4	How would you feel if someone sold your X's data without your permission? Why?
#	Behaviour Questions
1	How do you use your X to interact with others? Why?
2	Does your X allow you to set a password? Have you set a password? Why?
3	How much time have you spent reading your X's privacy policies? Why?
4	How much time have you spent configuring your X's privacy settings? Why?
5	How would you use your X if you wished to protect your privacy? Why?
#	Privacy Paradox Questions
1	Why do you think some people use devices which place their privacy at risk?
2	Why do you think some people use their devices in an unprivate way?
3	Why do you think some people claim to value privacy but still use devices which place their privacy at risk?

they might appreciate that unease is reasonable. This also sought to extract privacy opinions, if any were pre-existing.

We then gradually introduced the theme, asking how often a participant switches off their product. Some devices, particularly IoT technologies, possess environmental sensors. If an individual has a smart TV and never turns it off, they might be

exposing data [108]. In the final question, we targeted privacy more directly. We asked users who they believed could access their data. This sought to evaluate their degree of privacy awareness. If they recognise the risk, they demonstrate an element of knowledge. With the topic now introduced, we then solicited user concerns.

Concern. Rather than asking individuals whether they were concerned (due to response biases), we solicited reactions to contextualised questions. This scenario-based approach has proved successful in many previous privacy studies [2, 133, 232]. To select these incidents, we made use of three criteria: *relevance*, *feasibility* and *behaviour correspondence*. These are used throughout this thesis and are justified in Section 3.5. For the *feasibility* criterion, we considered whether the issue was possible in our *general* threat model (Section 3.4).

Based on these criteria, we selected four incidents. These comprised: *data deletion*, *data sharing*, *data monitoring* and *data selling*. The rationale for their inclusion is now discussed below.

Concern: Contextualised questions. *Data deletion* represents an intrusion into solitude/agency, one of Solove’s archetypal violations [352]. Since it infringes on a person’s property, it bears relevance to privacy. The incident is also feasible, particularly with the prevalence of hackers and malware. Ransomware has demonstrated the personal damage when files become inaccessible [172]. Finally, this issue could be partially mitigated through password usage. If a product is not locked, it might be easier for an adversary to access. Fortunately, passwords are a common feature to enhance protection. Based on the above considerations, this incident was deemed apt. When a person opposes the loss of data, they show concern for this information.

Data sharing has relevance to privacy, directly relating to Solove’s disclosure principle [352]. Sharing can also expose details to unintended parties, placing confidentiality at risk. This incident is highly feasible, since sharing takes place as a frequent event. IoT devices often synch with remote servers [60], transmitting data over long distances. Sharing is also common on the web, with this action supported by social networking sites. Finally, this issue can be limited by available techniques. Privacy/security settings are common to a range of products. By restricting these permissions, data disclosure can often be limited.

Data monitoring relates to the surveillance violation [352], since it concerns device observation. If a product is being tracked by another party, it might be an invasion of privacy. This incident has feasibility, with surveillance highlighted by the Snowden revelations [337]. More recently, the Facebook controversy showcased how some companies monitor user activity [75]. Many products communicate with their

vendor, whether to install updates or synchronise data. This provides further opportunities for activities to be tracked. With observation now being commonplace, some have accepted this monitoring [353]. Therefore, if a user rejects tracking, they demonstrate privacy concern. Finally, this risk might be mitigated through protective behaviour. Settings often allow such interactions to be restricted, as seen in the case of permissions. If users constrain configurations, they might limit data monitoring.

Data selling was the final incident, with this relating to the secondary use principle [352]. If information is provided for one purpose, and then used for another, this could be a privacy violation. Data selling is also feasible, with information regularly sold on online markets. Recent controversies have demonstrated how details are traded between brokers [239], with this asset aiding the digital economy [334]. Finally, there is a correspondence between the scenario and protective behaviour. Data can only be sold if it is acquired by a company. By restricting settings, the risk of such collection is reduced. Even if some information is obtained, users may be able to limit the quantity. If a person rejects this incident, they appear to show concern for their data.

Concern: Scores. After the interviews, the responses were transcribed from our audio recordings. Based on the replies to all four incidents, we formed the *concern score*. By exploring privacy from several angles, we provide a rich analysis of participant opinions. The coding process is described in detail in Section 6.3. This *concern score* represents one half of the Paradox.

Behaviour. We then explored behaviour; the second half of the Paradox. The section began with a general question, asking individuals how they use their devices. This was included to familiarise the participant with the new topic. Since the query was not grounded in a particular technique, it was not factored into the *behaviour score*.

Protective measures were also selected based on four criteria. Again, their use is justified in Section 3.5. They comprised: *simplicity*, *utility*, *applicability* and *concern correspondence*. Based on these criteria, we selected three privacy-protective features: configuring *passwords/PINs* (screen locks), reading *privacy policies* and adjusting *privacy settings*. The rationale for their selection is now discussed below.

Behaviour: Features. *Passwords* are familiar measures, used regularly by many individuals. They should be simple and understandable, fulfilling the simplicity criteria. Passwords protect personal data, reducing the risk of unauthorised access. Furthermore, screen locks have been recommended to mitigate theft risk [102]. Therefore, this technique possesses the necessary utility. Passwords/PINs are found on a diverse range of technologies, from desktops to smartwatches, laptops to smart appliances.

This suggests they have sound applicability. Finally, passwords can reduce the risk of unauthorised access, whether physical or virtual. This should limit the likelihood of several concern incidents. In particular, the security benefits should make data deletion less probable. Therefore, there is a correspondence between concerns and this approach.

Privacy policies are also well-known, found everywhere from websites to software licenses [235]. While the documents themselves can be opaque, their underlying purpose is simple. Now that the EU General Data Protection Regulation (GDPR) has been implemented, they should increase in accessibility. Policies also have clear utility, granting an individual greater privacy awareness. As previously discussed, this might encourage the adoption of protection [48]. Based on this knowledge, they could choose to reject a service or select a preferable alternative. These documents may also have a cumulative effect, supporting the use of passwords or privacy settings. The *applicability* criterion is fulfilled, as policies are supplied with a wide range of products. Again, the GDPR mandates greater transparency over data processing. Finally, invasive practices are often enumerated in these terms. Clauses will frequently explain how data can be shared, monitored or sold by a vendor. These issues align with three of our aforementioned incidents. Therefore, policies offer the information to respond to concerns.

Privacy settings, including permissions and sharing preferences, are found on a variety of products. While they often present an ‘all-or-nothing’ choice [35], they are a rare means to limit data collection. Individuals might have seen these menus, since they are found on apps and social media platforms. Furthermore, ‘settings’ should be a simple and understandable term. Providing utility, these tools allow data access to be restricted. Even if some information is still collected, many details tend to be optional. This has been found on both desktop systems [171] and mobile environments [288]. Most importantly, they provide protection against the concern incidents. If an individual wishes to limit sharing or monitoring, adjusting settings is a wise response. This has the added benefit of reducing opportunities for data selling. By selecting feasible protective measures, we avoid comparing concerns with impractical actions.

Behaviour: Scores. Based on whether the participant enabled a password, read their policies and adjusted their settings, we generated the *behaviour score*. This coding process is described in detail in Section 6.3. The score comprised the second half of the Paradox, and was compared against claimed concern.

Bias mitigation. It is appreciated that self-reporting can be prone to bias [134]. However, we took several approaches to minimise this issue. Firstly, we sought to

disguise our study’s privacy focus. Therefore, social desirability bias and demand characteristics should not have encouraged fabrication. Secondly, since compensation was allocated evenly, there was no financial incentive for misreporting. Thirdly, the behaviour in question was not sensitive or stigmatised. For this reason, participants should feel comfortable in responding. Finally, interviewee anonymity was emphasised and quotes were not attributed. Therefore, participants would not be embarrassed by the publication of results. Finally, the requirement for rationale may serve as disincentive. Since participants had to justify their claims, we thought they might be less likely to fabricate. These approaches should minimise the risk from self-reporting. As 77.5% of interviewees admitted to not reading their policies, it appears unlikely that behaviour was exaggerated.

Privacy Paradox questions. Our final three queries sought to explore justifications behind the Paradox. Once this rationale was understood, we could then develop approaches to mitigate the issue.

We considered asking participants why their behaviour differed from their concerns. However, we believed this approach would contribute to defensive or evasive responses. Therefore, our queries were phrased in terms of why ‘other people’ might act in this manner. Since their views would remain salient from previous responses, we expected answers to still correspond with their own rationale. 21/40 participants referred to themselves while responding. Since this might have been implicit in additional cases, we believe our approach was successful.

Paradox rationale was triangulated through three questions. The first asked why some people use devices which place their privacy at risk. We then queried why some individuals might use their products in an ‘unprivate’ manner. The term ‘unprivate’ was selected as it was more accessible than describing a lack of protection. Finally, we addressed the central question behind the Paradox: *Why do some people claim to value privacy but still use devices which place their privacy at risk?* With this question revealing the interview’s purpose, it was placed last to avoid order effects. Through dissecting each participant’s rationale, we explore a range of potential factors. This allows us to address our fourth subquestion¹, and supports the development of mitigative approaches.

¹SQ4: Which factors contribute to the Privacy Paradox in the IoT?

6.3 Results and Discussion

Participants

Category distribution. We conducted semi-structured interviews with 40 participants. 20 of these discussed their IoT device, while 20 outlined their ‘non-IoT’ product (as defined in Chapter 5). In terms of the IoT, 13 had wearables, 5 had smart appliances and 2 had home automation systems. For the ‘non-IoT’ group, 15 had laptops, 4 had tablets and 1 had a desktop computer.

We expect such a distribution to be due to two reasons. Firstly, as demonstrated in Chapter 5, some products are more popular than others. Secondly, if an individual owned several devices, they were asked to discuss one they used frequently. This sought to support informed responses to our interview questions. With wearables being worn on the wrist, they might be accessed more often than a tablet. Similarly, a laptop might be more convenient than a static desktop. Although the category sizes were far from equal, they represent the popularity of the assessed products. To compare the IoT and ‘non-IoT’ environments, we deemed it apt to consider this factor. The matter is discussed in our research critique (Section 9.3).

Demographics. Of these 40 participants, 60% were male and 40% were female. This closely corresponds with the 57%/43% split in our online survey (Chapter 5). The sample was quite young, with 45% between 26-35 and 32.5% between 18-25. This contributed to an estimated mean age of 31.6. This might have been due to the interviews being conducted during weekdays. We expect older individuals would be more likely to be in work or caring for children.

52.5% of our group possessed a Master’s qualification, with another 32.5% having a Bachelor’s degree. These proportions are likely due to the demographics of Oxford, and particularly those keen to participate in research. However, the estimated mean age (31.6) is much greater than the undergraduate average. This suggests that students did not dominate our sample.

Participants considered themselves to be technologically-literate ($\bar{x} = 4.0$), with 43% rating themselves 4/5. Those skilled in technology are often better-able to protect themselves [291]. This might place a minimum bound on Paradox prevalence. Interviews had a mean duration of 17.7 minutes, ranging from 11.5 to 28.9. All participants were asked the same questions, but some were slower and some were more open to elaboration.

Qualitative analyses

We transcribed our 40 interviews in their entirety, seeking to accurately record the comments. The doctoral student undertook ‘verbatim transcription’ [246], since they were familiarised with the data. Although this approach was time-consuming, it minimised the risk that relevant details were omitted [246].

Coding frame creation. We then undertook inductive analysis [323], as described in Section 3.3. Through using the same methodology, we conducted consistent explorations. As required, we first familiarised ourselves with the data. After undertaking this process for all 40 interviews, general themes began to emerge. For example, while some rejected policies due to complexity, others were deterred by the effort. Based on these themes, we created draft coding indices. Since our queries addressed privacy from several angles, we developed an index for each question. This approach supports the comparison of responses between participants [187].

We also noted ‘deviant cases’ [23], avoiding the temptation to subsume infrequent opinions. Once the indices were refined, they became our coding frames. One frame was produced for each of our 22 questions. Strict definitions were then added, seeking to make this process robust and replicable. The coding frames relevant to our current discussion can be found in Appendix A.2.

Coding approach. In answering a question, a participant might mention several topics. For example, Question 6 asked participants whether they had considered upgrading their device. Rather than including ‘Yes’ and ‘No’, we opted for ‘Reasons in favour’ and ‘Reasons against’. This allowed us to remain faithful to the diversity of responses. For behavioural questions we took a different approach. These responses indicated whether a technique was used or not. Therefore, top-level themes were categorised in an exclusive manner. As an example, Question 15 focused on device passwords. Since participants reported either using or avoiding these features, responses were coded as ‘Yes’ or ‘No’. However, individuals could still provide multiple subtheme reasons for their answer. Section 3.3 outlined how theme proportions could be reported in several ways. We primarily adopt the ‘comment’ technique, with the ‘participant’ approach signposted when used.

Qualitative validation. After we completed the coding, validation was undertaken to add further rigour. We considered using ‘multiple coding’, by having a second researcher analyse the transcripts [293]. However, this approach was deemed impractical for two reasons. Firstly, the research was undertaken by a single doctoral student. Therefore, we did not have an available resource at this point. Secondly,

with with forty 20-minute discussions, the time investment would have been significant. This also deterred us from a lengthy process of ‘repeated coding’. Since Chapter 8 addressed our central research question², the effort was considered better-spent on those interviews.

However, we wished to validate the appropriateness of our findings. To achieve this, we undertook respondent validation [56]. This process is outlined in detail in Section 3.3. Individuals were each sent their interview transcript. Beneath the answers, their relevant codes were listed clearly. They were asked to check both: a) the textual accuracy of the transcript; and b) the appropriateness of the assigned codes. Via email response, they could then request amendments on either grounds. As previously mentioned, these requests were not automatically honoured. They were balanced against our own academic judgement. Regardless of the outcome, interviewees were compensated with a £5 Amazon voucher. Since this required further cooperation after a lengthy interview, we believe it to be appropriate.

Validation results. After contacting our 40 participants, we received 36 responses. Individuals were sent two emails over two weeks, and therefore had ample opportunity to reply. Of the 36 respondents, 33 were pleased with both their transcript and the assigned themes. This accuracy rate of 91.7% suggests that we were successful in encapsulating opinions. Of the three requested amendments, none sought to ‘improve’ our interpretation of behaviour. This suggests that protective actions were not underrepresented.

Participant 6 wished their Question 3 code to be altered. They highlighted that “*the screen is fantastic*” referred to interface quality rather than monitor size. Since this owed to an ambiguous statement, the coding was updated. Participant 20 disliked the ‘Not interested in exercise’ code attached to their Question 7 response. This query did not explore their own rationale, but asked why other people might not purchase fitness trackers. Based on the interviewee’s feedback, we felt the ‘No need’ theme was more appropriate. Finally, Participant 24 identified a small transcription error in their Question 12 response. In the phrase “*because I have an urge to do it*”, the word ‘don’t’ was mistakenly omitted. However, since the respondent agreed with the assigned code, the meaning of the statement did not change. In this case, the transcription error was simply corrected.

Summary. Respondent validation supported the rigour of our qualitative analysis [56]. We also used triangulation to build confidence in our findings. Rather than

²Central question: *Can the Privacy Paradox be mitigated in the context of smartwatches?*

exploring the Paradox through a single query, we included three. Similarly, we assessed concerns through four contextualised questions. Furthermore, behaviour was evaluated based on three actions. Through using triangulation, we can corroborate our findings.

Score generation

Concern scores. To assess concern and behaviour, we required comparable metrics. However, we were concerned that qualitative comparisons might be excessively subjective. For example, if a person justifies their lack of protection, which factors could correspond with concern? We therefore decided to compare quantitative metrics. These were calculated in a consistent manner from participant responses. These replies had already been rigorously coded and validated by respondents. When individuals expressed a large number of objections, they received a high *concern score*. Similarly, if protective measures were all used, a high *behaviour score* was assigned. The process for these calculations is outlined below.

Both metrics encapsulated multiple responses and were assessed on a five-point scale. In this manner, they provided an overview for each participant. If the scores differed greatly, a disparity could be identified. Justifications could then be extracted from the rich data. Without this approach, it would be challenging to compare Paradox prevalence. We reflect on the process in our critique (Section 9.3).

Concerns were gauged based on responses to the four contextualised questions. For each query, concern intensity was assessed in the following manner. As shown below in the ‘data deletion’ coding frame (Table 6.2), comments were grouped into ‘Concerns’ and ‘In mitigation’. While definitions are omitted for brevity, the full tables can be found in Appendix A.2.

Since interviews are conversational, participants often responded with several comments [307]. If all a participant’s remarks were concerns, their reaction was marked as ‘Concerned’. This appears appropriate, as they did not express any reasons in mitigation. In contrast, if none of the comments were concerns, they were marked as ‘Not Concerned’. We believe this also apt, since individuals were given an opportunity to report their worries. Indeed, respondents often spent several minutes outlining their many considerations. If a negative incident is greeted with no apprehension, then we believe the label is fair. Finally, if the participant expressed both concerns and mitigations, they were classified as ‘Slightly Concerned’. This was deemed apt, as the individual appreciates both the threats and the limitations. While we considered additional divisions, we thought the differences would become marginal.

Table 6.2: Example Coding Frame: Data Deletion Incident

Theme	Subtheme	Subsubtheme
Concerns	Data	Valued
		Reliant
	Recovery	Effort to recover
		Cannot fully recover
	Principle	Principles
		Unconsented
		Invested in device
		Feels like theft
		Security concerns
		Lost functionality
		Inconvenience
In mitigation	Contingent	Depends what data
		Depends how/why deleted
		Innocuous data
		Backed-up data
		Simple to replace
		No long-term requirement
		Other matters are more important

Our process was undertaken for all four contextualised questions, with each having ‘Concerns’ and ‘In mitigation’ themes. This enabled a holistic analysis of privacy concern. Once complete, we began generating metrics. For each question, 0 was assigned to ‘Not Concerned’, 1 to ‘Slightly Concerned’ and 2 to ‘Concerned’. This resulted in cumulative metrics from 0 (low) to 8 (high). Appropriately, an individual expressing widespread concerns would receive a higher total than one with varied responses.

Finally, to enable comparison with behaviour, metrics were scaled to a score between 1 and 5. Our approach is illustrated in Figure 6.1. The technique sought to discourage indecision by granting the neutral 3/5 a smaller region. This was undertaken since we wished to differentiate between high and low concerns. However, if balanced responses were given for all four scenarios, the neutral score was assigned. The final metric, known as the *concern score*, allowed opinions to be compared across our 40 participants. If an individual expresses many concerns to many incidents, they appropriately receive a large score. However, if their opposition is moderated by other considerations, they get a lower metric.

Behaviour scores. We followed a similar approach to analyse behaviour. The metric was calculated based on the usage of the three protective techniques. Since participants either did or did not use these tools, only a single top-level theme was

Scenario	Reaction	Metric	Total		Score
Q10 Data Deletion	Concerned	2	8	}	5
	Slightly Concerned	1	7		
	Not Concerned	0			
-----			6	}	4
Q11 Data Sharing	Concerned	2			
	Slightly Concerned	1	5		
	Not Concerned	0			
-----			4	}	3
Q12 Data Monitoring	Concerned	2	Σ	}	2
	Slightly Concerned	1			
	Not Concerned	0		2	
-----			1	}	1
Q13 Data Selling	Concerned	2			
	Slightly Concerned	1			
	Not Concerned	0	0		

Figure 6.1: Qualitative Concern Scaling

applicable. Therefore, participant behaviour was easier to categorise. An example coding frame, that of password usage can be found in Table 6.3.

Table 6.3: Example Coding Frame: Password Usage

Theme	Subtheme	Subsubtheme
Yes	Security	General protection
		Physical access
		Loss/theft
		Easy/quick
		Parental controls
		Save battery
		Habit
		Requirement
		Duty to protect work device
		Separate space
		Sensible
		Assisted by another
	No	
		Not supported
		Too much effort
		Little perceived risk
		Data considered innocuous
		Have not investigated
		Only for online account

To comprise the *behaviour scores*, points were assigned based on protection. All participants began with a single mark out of five. If they indicated that they used

a password, they were given 1 additional mark. They were assigned another mark if they read their privacy policy in detail. Superficial skimming was not rewarded, since it would not provide sufficient awareness. Finally, if participants configured their settings, they were given 2 extra marks. This technique was weighted heavily since it offers protection against multiple incidents. As previously outlined, permissions could limit data sharing, monitoring and selling. The total, deemed the *behaviour score*, ranged between 1/5 (low) and 5/5 (high). A person who used all protection would appropriately receive 5/5. In contrast, if only policies were considered, 2/5 would be assigned. We deemed this metric to be apt, since it illustrates the usage of protection. Through this scale, we compared behaviour between our participants.

General responses

These questions sought to both acclimatise our interviewees and disguise privacy. We now outline responses and compare our two groups (IoT and ‘non-IoT’). For brevity, we only discuss those queries with findings relevant to the thesis.

What do you like? When naming their most popular feature, our groups greatly differed. Whereas the ‘non-IoT’ group valued the form factor (39.2%) and usability (25.5%), IoT owners cared about applications (54.2%). In the latter case, sensors and exercise tracking were particularly enjoyed. This echoes the focus on functionality found in the Chapter 5 survey. Privacy was not cited by a single interviewee. Since tracking is in tension with this factor, features appear more appreciated.

What do you dislike? Across both groups, the largest objections were to form factors (34.9%) and poor usability (20.6%). However, almost two-thirds of these criticisms came from ‘non-IoT’ owners. This was frequently due to poor casing, though heaviness was also an issue. As before, privacy was not mentioned. It appears that the matter often fails to be a salient consideration.

Why might others not purchase? While individuals might not have privacy concerns themselves, they appreciate that such views are valid. When answering this question, the most popular justification was that the product was unnecessary (22.9%). However, two IoT owners identified issues related to privacy. The first mentioned that smart TVs are suspected of eavesdropping. The second highlighted that people might be paranoid of Fitbit tracking. While we were pleased that privacy was mentioned, these issues were not raised in the ‘dislike’ question. Therefore, users might possess latent concerns which do not influence behaviour. Representative quotes, along with the participant number and group, are presented below:

“A lot of TVs now have voice control, I think people are a bit concerned about the privacy aspects of that” (#13, IoT).

“I think there are a lot of paranoid people who don’t really particularly like to wear any tech” (#18, IoT).

How often is it turned off? While this question did not ostensibly concern privacy, IoT products can collect data surreptitiously [108]. If such a device is left on constantly, information might be collected.

Usage patterns differed considerably based on the technology. Whereas 60% of ‘non-IoT’ owners switched off daily, only 20% of IoT participants acted in the same manner. Most of the former group sought to save battery/energy (58.8%), though some acted out of habit. Worryingly, 66.7% of IoT users switched off rarely or never. This was primarily to either provide features or enable data collection. It appears that many are aware of monitoring, but desire it to enhance functionality. If these individuals still claim concern, this might provide rationale for the Paradox.

“Just because I like to have the most data available to look at for myself, rather than having gaps.” (#35, IoT).

Who can access data? This question began to introduce the topic of privacy. If an individual lacks privacy awareness, they might be unable to perceive risks. Furthermore, they might lack the skills to protect themselves.

Of concern, ‘non-IoT’ owners were able to identify more issues (56.6%) than IoT users (43.4%). This was counter-intuitive, since many smart devices function through collecting data [300]. When considering those participants who were unsure, 8/11 (72.7%) were IoT owners. These users might deem the matter too complex. Seven individuals claimed to have not considered access. Worryingly, all seven owned IoT products. If users have little awareness, their information might be placed at risk.

“Essentially I have no idea who has access to my data.” (#25, IoT).

Privacy concerns

We move forward to consider the privacy concerns expressed by the participants. Our questions solicited reactions to data deletion, sharing, surveillance and selling. By comparing user responses, we can explore which devices provoke the most concern.

Deletion. Across both groups, 92.5% expressed some concern over deletion. Since 45% were strictly opposed to the act, individuals appear to feel some sense of ownership. With such opposition, it would appear rational for protective measures to be used. While the concern did not differ between groups, the justifications did. ‘Non-IoT’ users opposed the principle of violation (43.6%), particularly disliking the lack of consent. Many also thought that their data was valuable (17.1%), possibly because it included important documents. In contrast, IoT concerns were frequently due to lost functionality (16.7%). If a person’s Fitbit data is deleted, this limits the benefits of exercise tracking. However, many doubted sensitivity, with half the mitigative comments mentioning this. As suggested by earlier responses, it appears IoT users place great value on functionality.

“Because I probably would lose things that I need for work.” (#11, Non-IoT).

“I would be pretty frustrated, because I use it as a training diary.” (#35, IoT).

Sharing. This practice was also opposed, with 90% of participants expressing at least some concern. The groups were united, with over 80% of both conveying this emotion. Despite the popularity of online disclosure [4], individuals appear to desire some control. To receive this capability, they could adjust their device’s settings.

Only 37.5% of participants showed strict opposition, implying it was not as unpopular as data deletion. The most contentious point was the principle of invasion, mentioned in 59.6% of comments. Whereas some might voluntarily share data, they oppose disclosure without their consent. IoT owners were concerned about how their information would be processed (26.7%). Several feared that GPS could be used to track geographical locations. Since GPS tends to be optional, this risk could be reduced by adjusting settings.

“I think that if information belongs to an individual then they have authority over what happens to that information.” (#26, Non-IoT).

“The Fitbit data is gradually saying too much about you. It says days when you’re active, it could tell where you are cycling.” (#4, IoT).

Monitoring. When considering surveillance, 90% of both groups showed opposition. 50% were strictly opposed to the practice, more than in the previous two questions. This suggests that many dislike tracking, even in our increasingly-surveilled society. This conflicts with the functionality of many products, particularly smart

devices [108]. ‘Non-IoT’ participants found monitoring uncomfortable and thought it might influence their behaviour (34.6%). In mitigation, some accepted that surveillance was a social norm (68.8%) while others had ‘nothing to hide’ (12.5%).

The rationale appeared to differ slightly for smart device owners. Although targeted advertising was disliked, tracking and security provoked the most concern (18.5%). Whereas a laptop might only disclose your IP address, a smartwatch might identify your location. With IoT users opposing such monitoring, one might expect them to adjust their settings.

“So I think it would probably change your behaviour definitely and again it wouldn’t be particularly pleasant.” (#3, Non-IoT).

“I’d feel like, like someone would maybe be stalking me which would be a bit unnerving.” (#35, IoT).

Selling. This practice faced strong condemnation from both groups. 92.5% disliked selling, with 62.5% expressing strict objections. At least 60% in both samples were purely opposed, suggesting this incident provoked the most concern. This is unsurprising, as individuals should reject being exploited for profits. Nevertheless, the views are at odds with the digital economy [334]. If participants are so opposed to this practice, one might expect them to guard their data.

Both groups took exception to selling, but for very different reasons. ‘Non-IoT’ owners disliked that another party was making a profit (26.3%). In contrast, IoT users wanted to receive money from the transactions (22.2%). Since they were not opposing the access itself, they might consider their data to be less valuable.

“Well that’s the kinda same as sharing actually, just slightly worse as they’re actually making money out of it.” (#17, Non-IoT).

“I would also be angry because I should get part of the money.” (#36, IoT).

Concern scores. Based on the responses, owners appeared concerned about their devices. However, to assess participants’ views, we calculated the *concern scores*. 82.5% received a metric of either 4/5 or 5/5. In contrast, only 10% got a score beneath 3/5. We continued by analysing our two groups. 75% of IoT owners received over 3/5, contributing to a mean score of 3.95/5. The ‘non-IoT’ group also showed opposition, resulting in a mean of 4.15. This suggests that most participants expressed privacy concerns. Since owners appeared to oppose violations, one might expect them to protect their data.

Privacy behaviour

To complement our evaluation of concerns, we now move on to consider behaviour. This was gauged on whether participants enabled passwords, read their privacy policies and configured their settings. If one studies terms and adjusts configurations, they might be better prepared than one who abstains.

Passwords. Password protection was far from perfect, with only 55% of participants securing their devices. However, this percentage was skewed due to the inaction of smart device owners. 95% of ‘non-IoT’ users set a password, which corresponds well with their strong concerns. In contrast, only 15% of IoT owners did the same. Based on this dichotomy in behaviour, we found passwords were used significantly less often on smart devices ($X^2(1) = 25.859, p < 0.001, d = 2.705$). This ‘huge’ effect size [333] suggests a considerable contrast in behaviour. If passwords are neglected, users miss an excellent opportunity to protect their data.

‘Non-IoT’ owners used passwords for several reasons. Many appreciated the security protection (21.4%) and the restriction of physical access (19.0%). Some also enabled this feature out of habit (9.5%), regarding configuration as part of their routine. IoT users appeared to view passwords in a different manner. They disliked the degree of effort required (23.1%), as their products often rely on brief interactions. Indeed, some devices might lack interfaces suited to password entry [183]. Other owners failed to investigate the features (19.2%) or doubted their device’s sensitivity (23.1%). While data can often appear innocuous, inference techniques may reveal private information [106].

“If anyone would break into my home, I wouldn’t want them to access my computer” (#12, Non-IoT).

“I just want to swipe it. It just takes too much time to get in there” (#21, IoT).

Policies. These documents provide a rare approach to learn about data procedures. Since our participants appeared concerned about privacy, one might expect them to take this opportunity. However, only 22.5% of our sample took time to read the policies. While we do not claim to be representative of the public, this low percentage is supported by existing work [155].

As before, behaviour differed greatly between our two groups. Whereas 40% of ‘non-IoT’ owners studied the documents, only 5% of others acted the same. Based on this, smart device users were significantly less likely to read their policies ($X^2(1) = 7.025, p = 0.02, d = 0.923$). This ‘large’ effect size further emphasises the contrast.

‘Non-IoT’ owners took action to improve their knowledge (28.6%). Some checked for egregious clauses (21.4%), while others found the documents quick to read (14.3%). In contrast, IoT users were deterred by the excessive length (11.8%) and the time required (14.7%). Many also commented that they were too eager to use functionality (14.7%). When individuals become preoccupied with features, they might forget their reservations [103]. Regardless of the rationale, IoT behaviour differed from concerns.

“I tend to take a reasonable glance and see if there’s any ridiculous stand out clauses” (#11, Non-IoT).

“I think I was more in a hurry to get it out of the box and set up and start using it” (#40, IoT).

Settings. Configuration appeared common, with 70% of our sample adjusting their options. However, as before, behaviour differed greatly between groups. While 95% of ‘non-IoT’ owners took action, only 45% of others did the same. Since default configurations tend to be permissive [59], users might place their privacy at risk.

When comparing groups, we found IoT owners were significantly less likely to adjust their settings ($X^2(1) = 11.905$, $p = 0.001$, $d = 1.301$). Once again, this ‘very large’ effect size [333] suggests a disparity in protection.

Rationale differed greatly between the two groups. When ‘non-IoT’ users took action, it was justified on improving privacy (23.1%) and security (23.1%). Many also enjoyed tinkering with their system to customise the experience (15.4%). IoT users appeared less interested, considered functionality to be more exciting (18.8%). If owners are preoccupied with features, they might quickly forget their concerns. These users also perceived little risk (12.5%) and doubted their data sensitivity (12.5%). While their judgements may be accurate, such decisions require a degree of knowledge. Since a 2017 poll found only 8% understood sharing practices [99], this knowledge might be rare. If true, low awareness might contribute to lax behaviour.

“When I first set up anything I tend to like to go through and adjust the dials” (#7, Non-IoT).

“I just want to explore the functions and interesting bits not the privacy bit, privacy is the boring bit” (#36, IoT).

Behaviour scores. To compare actions between participants, we calculated the *behaviour scores*. 52.5% received a metric above 3/5, with these scores indicating protective behaviour. This was compared to 30% who got a metric beneath 3/5. It appears that while most use some protection, many remain at risk.

In the ‘non-IoT’ group, 90% received a score of 4/5 or above. This resulted in a mean of 4.25/5. The metric implies that these users take effort to protect their devices. However, only 15% of IoT owners received such high scores. Indeed, 55% of their group were assigned 2/5 or below. Their mean of 2.1/5 demonstrates how infrequently protection is used. With this metric appearing much lower than *concern scores* ($\bar{x} = 3.95/5$), a disparity seems to be likely.

H1. We now turn to our first hypothesis. It asserted that the IoT mean *behaviour score* will be significantly less than that for ‘non-IoT’ products. If accepted, it might imply that smart devices constrain protective behaviour. Through a Mann-Whitney U test, we found that the difference was indeed significant ($U = 30.5$, $p < 0.001$, $d = 2.105$). The ‘huge’ effect size further suggests a gulf in behaviour. Based on the aforementioned justifications, this appears due to two factors: a preoccupation with functionality and a lack of awareness. As will be highlighted, these reasons were also prevalent when discussing the Paradox.

Categories. To determine which devices pose the greatest risk, we compared behaviour between our categories. If a product is rarely protected, its data might be vulnerable to attack. Since the online survey (Chapter 5) suggested wearables lack usability and familiarity, we expected them to present the greatest issues [230, 317].

A Kruskal-Wallis test showed that *behaviour scores* differed significantly based on device ($X^2(5) = 26.162$, $p < 0.001$). In order to determine the sequence, pairwise analyses were then undertaken ($\alpha = 0.003$). We received two significant results, with other comparisons hampered by low quantities of devices. Smart appliance behaviour was significantly less-protective than that on laptops ($U = 0$, $p < 0.001$, $d = 2.148$). Wearable protection was also less used than tools on laptops ($U = 7$, $p < 0.001$, $d = 3.211$). This effect size was even greater, emphasising the infrequency of action. Wearables contributed to 13/20 IoT products, and hence received a detailed analysis. Since they might constrain protection, they appear a good candidate for investigation.

The Privacy Paradox

Our third research subquestion considers Paradox prevalence in IoT and other environments. An absence of protection does not alone indicate that the issue is present. To explore this subquestion, and our second hypothesis, we compare *concern scores*

with *behaviour scores*. Comparisons were undertaken on a per-participant basis, rather than considering the aggregate metrics. This sought to ensure that concerns and behaviour were contextualised around the same product. We defined a two-point gap as denoting a *disparity*. Such occurrences suggest that the Paradox is present.

Prevalence. Based on this definition, 12/40 (30%) participants were prone to the issue. As previously stated, our study does not claim to be representative of the general public. However, it suggests that a subset of individuals might be susceptible. This result corroborates the confirmation of the disparity in our street survey. Of these 12 participants, 11 (91.6%) owned IoT devices. Based on this division, there appears to be an issue in smart device environments.

H2. Our second hypothesis posited that the *disparity* would be significantly more frequent for IoT devices than for ‘non-IoT’ products. Based on the results of a Chi-Squared test, we accepted this conjecture ($X^2(1) = 11.905$, $p = 0.001$, $d = 1.302$). Even if we selected a 3-point *disparity* definition, 9/40 (22.5%) would still be prone. The resulting 8-1 division would also remain statistically significant ($X^2(1) = 7.025$, $p = 0.02$, $d = 0.923$). In accepting the hypothesis, the IoT appears more prone to the Paradox. This addresses our third research subquestion. If smart devices do exacerbate the issue, their proliferation might place privacy at risk.

Paradox distribution. Below in Figure 6.2, we directly plot *concern scores* against *behaviour scores*. We represent participants with individual points, with red denoting IoT users and blue denoting ‘non-IoT’ owners. The heatmap is divided into three regions: commensurate behaviour (yellow), stronger-than-expected behaviour (green) and the Privacy Paradox (pink). Those 25 in the yellow area appear to act roughly consistently with their concerns. Individuals in the green region (3) take more action than their concerns suggest, possibly due to routine or expertise. The Paradox is represented by the pink area; where concern is accompanied by little protection. Based on the distribution of points, the IoT appears prone to the issue.

Categories. We then analysed the matter through our categories. By identifying devices which cause risk, we could select an environment for further study. Since many of our products had a frequency lower than 5, the Chi-Squared test was inappropriate [137]. Therefore, we opted for Fisher’s Exact Test instead [137]. Through this, we found that Paradox prevalence differed significantly across the devices ($p = 0.003$). The two products which contributed the most cases were smart appliances (3/5 users) and wearables (7/13 users). The latter category was expected, since it received poor evaluations in the Chapter 5 survey. Although the proportion is larger for appliances,

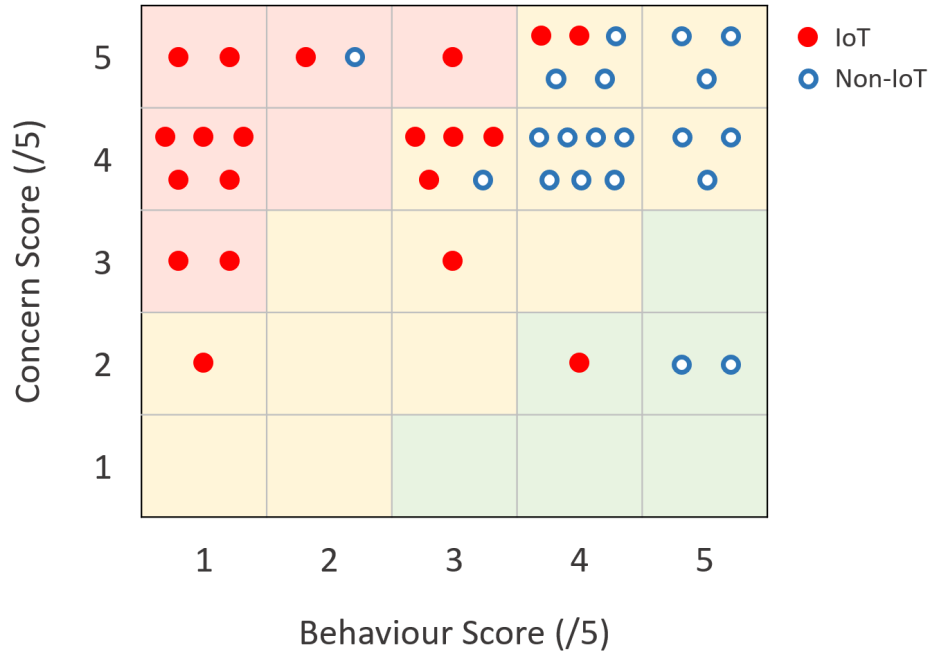


Figure 6.2: Heatmap: Concern Scores vs Behaviour Scores

wearables comprised over half of the Paradox incidents. Therefore, we believe that they present the greatest challenge to protective behaviour.

Contributory factors. Through accepting our second hypothesis, IoT owners appear more prone to the Paradox. However, to mitigate this disparity, we must understand the contributory factors. This allows us to address our fourth research subquestion. To dissect the decision-making process, we analysed our final three queries. They explored why ‘risky’ devices are used, why protection is neglected and why the Paradox exists, respectively.

Lack of awareness was the most frequent Paradox justification. It was mentioned 30.8% of times for the first question, 34.4% for the second and 19.3% for the third. In all three cases, it was the most popular reason. In terms of participants, it was cited by 28/40, 26/40 and 16/40, respectively. Most importantly, it was the most-common IoT factor for the final two questions. 4/12 (33.3%) Paradox-prone individuals used this point to justify the disparity.

Some thought that people did not recognise the threat. In this case, it is understandable that protective tools were not investigated. Others believed that even if protection is sought, most do not know of the techniques. It matters little that risk is recognised if remediation is not understood. This is supported by Büchi et al. [72], who found that Internet skills are more influential than concerns. It also corroborates Chapter 4, which found that those with expertise use more protection. Again, this

highlights the two components of privacy awareness: threat and response. Therefore, to mitigate the Paradox, we could seek to increase privacy awareness. Representative quotes are shown below, along with participant number, group and disparity status:

“I guess it’s just they don’t always know what people can access and what they can do with that information. So people like me who don’t read the terms and conditions and the privacy settings.” (#28, IoT, Disparity).

Convenience benefits were also found to be a popular justification. Products might need data for functionality, and therefore require settings to be open. This factor was cited 23.1% of the time to justify ‘risky’ devices, 11.5% to explain a lack of protection, and 13.6% when considering the Paradox. In terms of participants, it was mentioned by 23/40, 11/40 and 12/40, respectively. It was also the most-popular IoT justification for the first question. Similar to a lack of awareness, 4/12 (33.3%) Paradox-prone individuals cited this reason.

This is supported by the Paradox work of Beuker [55], who found that perceived benefits encourage data disclosure. Exact reasoning differed by participant, but many described a balance between privacy and convenience. Whereas they valued the principle, data would be sacrificed to receive functionality. This exchange may be perceived as low-risk by the individuals. However, without privacy understanding, threats cannot be evaluated accurately [350]. Furthermore, those with less knowledge tend to overestimate the advantages [157]. Therefore, a lack of awareness might have further influence. Risky practices were also justified on short-term necessity. For example, even if public Wi-Fi is insecure, it might be used to send an urgent email. Individuals may aspire to privacy, but sacrifice their data through practical necessity.

“I think for convenience, I think convenience often trumps security considerations, at least in my case” (#25, IoT, Disparity).

Lack of risk perception was a common justification. The factor was used to justify ‘risky’ devices (14.4%), inadequate protection (7.3%) and the Paradox (13.6%). In terms of participants, it was cited by 12/40, 6/40 and 10/40, respectively. As before, 4/12 (33.3%) Paradox-prone individuals used this to justify the disparity.

Some users considered their data to be innocuous, and therefore not worthy of attention. Others trusted they were safe due to vendor protection. Still others appreciated that attacks were possible, but questioned their severity. These are all sensible considerations, and some judgements might be accurate. However, as suggested above, the public often lack privacy awareness [49]. Therefore, they may be

unable to ascertain their degree of risk [350]. Through increasing knowledge of both threats and protection, individuals might learn how to guard themselves.

“I don’t do it on purpose...you just assume that it’s new and it’s all protected”
(#21, IoT, Disparity).

Based on these findings, a lack of awareness appears to be influential. We accept that this justification might occasionally be used as an excuse. Some might not care about privacy, but claim ignorance through social norms [138]. While possible, many participants demonstrated little knowledge of data access. Furthermore, the influence of privacy literacy is supported by much research [48, 48]. Therefore, we believe lack of awareness to be an important factor.

To encourage protection, we will seek to impart privacy knowledge in our subsequent studies. As concerns are already strong, we do not expect them to further increase. Therefore, an educational approach might mitigate the prevalence of the Paradox. While awareness campaigns might appear an obvious solution, they are criticised for lacking efficacy [38]. This issue is highlighted within the Literature Review (Section 2.5). Individuals must be incentivised to act, otherwise they will not invest effort [38]. To adopt new behaviour, users must also be confident they can put knowledge into practice. Since campaigns lack efficacy, we explore interactive approaches to encourage protection. Even if individuals then avoid these tools, at least we have supported informed decisions.

6.4 Implications

Awareness. We found IoT owners used less privacy protection than other users. They took less action in setting passwords, reading policies and adjusting their settings. Since smart devices can collect large amounts of data [300], such behaviour might place privacy at risk. Qualitative justifications suggest this was often due to a lack of awareness. This factor was the most prevalent reason given by the IoT owners. It occurred because users did not know of the risk and/or did not know of the response. If individuals do not feel at threat, they have little reason to seek protection.

Paradox. Since concerns were common in both groups, the Paradox was more prevalent in the IoT. This corroborates our online survey (Chapter 5), which found concerns were high but usability and familiarity was poor. It also supports our street

study (Chapter 4), by demonstrating that the Paradox is common. If IoT devices are prone to this issue, their design might limit protection.

Wearables. Since wearables fared worst in our previous study, we expected them to be particularly susceptible. This was found to be true, with over half of their owners displaying the disparity. Based on these findings, smartwatches might pose a risk as the devices spread. This echoes the views of Barth and de Jong [47], who opined that the Paradox might be exacerbated by mobile computing. Such results encourage the analysis of wearable devices.

Next steps. This chapter addressed the third subquestion by comparing Paradox prevalence across our two groups. With the disparity appearing greater in the IoT, we are encouraged to explore novel environments. It also answered the fourth, since it provided an analysis of contributory factors. This supports the development of mitigative approaches, which we describe in the following chapters. Furthermore, we scope our focus to smartwatches primarily based on these interviews. In this manner, our subsequent research builds on this work.

To reduce the disparity, we seek to increase the privacy awareness of smartwatch users. However, public campaigns merely highlight a problem, rather than encourage new actions [38]. For behaviour change to be successful, users must be incentivised and given chances to refine their knowledge [332]. Therefore, we designed an interactive approach to encourage protection. Although action does not itself mitigate the Paradox, we do not expect concerns to greatly increase. Indeed, if participants possess protective knowledge, their fears might be moderated. As the Paradox has been confirmed and factors have been identified, we now turn to addressing the issue.

Chapter 7

Can we mitigate the short-term prevalence of the Privacy Paradox on smartwatches?

7.1 Introduction

Background

In the preceding chapters, we have demonstrated that the Paradox is more prevalent in the IoT. Wearable devices were found to be particularly susceptible, contributing to over half of the issues. This is unsurprising, as mobile computing has been claimed to support irrational decision-making [47]. With wearable users placed under the greatest risk, we wish to protect individuals with smartwatches.

In addition to gauging the prevalence of the Paradox, we have also sought to understand it. Through our interviews, the most common IoT justification was a lack of awareness. Although it might occasionally be used as an excuse, protection is challenging without knowledge. Indeed, Internet skills have been found more indicative of actions than concerns [72]. But while campaigns can highlight a topic, they frequently fail to change behaviour [38]. Such a result would challenge the mitigation of the Paradox. While we could seek to reduce concerns, this was deemed dishonest and unethical. Individuals do face a risk to their data and this should not be disguised. Therefore, we explored other approaches to encourage protection.

The literature recommends complementing awareness with both education and training [332]. Since we wish to highlight risks and teach protection, ‘serious games’ appeared wise [324]. We did consider alternative approaches, including campaigns, nudging and Privacy by Design (PbD). All these techniques are outlined in Section 2.5. However, as described in that section, they were deemed insufficient for several

reasons. Firstly, campaigns are generic, non-interactive and frequently ineffective [38]. Secondly, while nudging can influence behaviour, actions revert when the interface is changed [70]. Finally, although PbD is beneficial when planning systems, smartwatch settings are already implemented [355].

In contrast, educational games are successful at imparting knowledge [101]. Indeed, they have been found more effective for learning and retention than traditional instruction [394]. Games are interactive and do not require complete control of the ecosystem. Furthermore, through incentives and positive reinforcement, persuasion can persist after gameplay [101]. Gaming scenarios have been previously proposed as a means to teach privacy [375]. Games have also proved successful in a range of cyber security environments [199,316,341]. Due to these many factors, we designed and evaluated an educational game.

Motivation

Privacy educational games are rare [363], with none developed in the context of smartwatches. Due to the constraints of this interface (e.g., small screen, few buttons) [90], there was a risk of an ineffective implementation. To minimise this probability, our designs were informed by Psychology [38,150] and Learning Science [248,312] principles. We also developed the application following best practices [26,312]. However, to support an effective game, prototyping appeared appropriate. Therefore, we first design and develop a prototype game. This application is hosted online, supporting its trialling with a large group of smartwatch owners. While the watch interface is simulated, such approaches are commonly in privacy research [206,315,383]. Indeed, in Jackson and Wang's [194] Paradox study, they simulate their interface. Through this large-scale analysis, we can collect behavioural results and participant feedback. Based on these findings, we can then refine the final application.

To assess our prototype, we solicited the concerns and behaviour of smartwatch owners. This established a baseline of how individuals perceive their devices. After these responses were captured, users interacted with our online game. This provided education and allowed owners to practice new behaviour. Finally, their actions were re-evaluated in follow-up interactions. This enabled a detailed analysis of whether our prototype was persuasive.

We wished to conduct our analyses in a scientifically-rigorous manner. Therefore, we divided our participants into a control group and a treatment group. Both groups answered the same initial and final questionnaires. However, the treatment group also played the online game. If behaviour becomes more privacy-protective, this

might suggest the prototype was persuasive. If the control group do not differ from pretest to posttest, then questionnaires should have introduced little bias. Through comparisons of behaviour, our application can be fairly assessed.

In this study, we seek to mitigate the Paradox over the short term. This directly aligns with our fifth subquestion¹. It is difficult to address the Paradox, with only one previous work having success [194]. This was achieved within a single session, rather than over a period of time. Furthermore, behaviour change tends to be a challenging task. Therefore, we felt it would be unwise to directly attempt a medium-term mitigation. It was more prudent to first conduct a smaller experiment. If the Paradox is reduced over several days, a greater period can then be explored. Since our final app is informed by this prototype, it should have a greater chance of persuasion.

7.2 Methodology

Hypotheses

Metrics. We begin by outlining the four metrics we studied. These comprised: the initial and final *concern scores*, and the initial and final *behaviour scores*.

Before we can explore whether perceptions have changed, we must establish a baseline. If later responses have adjusted, it might be due to the game’s influence. Since privacy is contextual [278], it is challenging to analyse in a fair manner. Therefore, as before, we make use of contextualised concern questions. To select them, the consistent criteria were used from Section 3.5. However, the *general* model was too broad for our scoped focus. Therefore, we adopted the *smartwatch* model from Section 3.4. By considering feasible issues, concerns could be assessed appropriately.

Before participants learned of a game, they completed an initial questionnaire. In that form, we solicited reactions to three privacy incidents. The approach was similar to the previous study, with the process outlined later in this section. Respondents indicated their concern on a five-point Likert Scale. This ranged from 1/5 (low) to 5/5 (high). By calculating a mean across the scenarios, we produced the initial *concern score*. Since responses were aggregated on a single theme, this approach was deemed appropriate [58, 78, 280]. If a person was opposed to several incidents, their score was higher than one with mixed views.

To receive a baseline, behaviour was also solicited in the initial questionnaire. We explored actions through three simple questions. Each query requested the usage

¹SQ5: *Can we mitigate the short-term prevalence of the Privacy Paradox on smartwatches?*

frequency of a protective technique. The selection criteria were used from Section 3.5, with the process also outlined later in this section. The techniques sought to mitigate the risks from the contextualised questions. Therefore, if person expressed strong concern, it would be rational for them to use protection. Responses were again supplied on a five-point Likert Scale. This scale ranged from 1/5 (infrequent) to 5/5 (frequent). We calculated another mean, generating the initial *behaviour score*. If an individual had consistent protection, their score was greater than one taking occasional action.

After completing the initial questionnaire, our sample was divided into two halves. These comprised a treatment group and a control group. The former played the prototype game, while the latter undertook no tasks. We selected this approach since it can mitigate confounding variables [344]. Respondents were unaware of their group membership and had no knowledge of the alternative condition.

Once the gameplay stage was finished, all participants completed a final questionnaire. The first section was identical to the initial form, enabling a fair comparison. Therefore, we received updated responses for concerns and behaviour. If knowledge moderated fears, they might report less concern. Similarly, if the prototype encouraged protection, their actions might differ. Through this, we received the final *concern scores* and *behaviour scores*.

H1. Our interviews suggested that a lack (low degree) of awareness contributes to the Paradox. This might be because individuals do not perceive any risk. Even if they do, they cannot act without a degree of knowledge. Since our prototype both highlights privacy and teaches techniques, we believe protective behaviour will be encouraged. Based on this belief, we formed our first hypothesis. We assert that the final mean *behaviour score* in the treatment group will be significantly higher than its initial mean score. If so, the game might have encouraged protective behaviour.

H2. To mitigate confounding variables, the control group do not play the game [344]. They complete no tasks between the initial questionnaire and the final questionnaire. Because of this, we expect their responses to be stable. However, if behaviour changed significantly, another factor may have influence. This might be the survey questions or the structure of the methodology. We expect this to be unlikely, since our questionnaire is designed to disguise the topic of privacy. Therefore, we posit that the final mean *behaviour score* in the control group will *not* be significantly higher than the initial mean score. Null hypothesis significance testing is a matter of continued debate [277]. However, we do not test for equivalence; we only care that differences

do not reach the threshold. If H1 and H2 are both accepted, then any behaviour change may be due to the prototype.

H3. An increase in protection does not itself mitigate the Paradox. When evaluating the issue, we compare concerns against behaviour. Since all questions were contextualised around an owned smartwatch, these comparisons can be undertaken fairly. Both *concern scores* and *behaviour scores* were rated on a five-point scale. These scores present a quantitative overview of each participant's responses. As before, we define the *disparity* as a two-point gap. In such cases, concern might be strong but protection used occasionally. Alternatively, opinions might be varied but action is never taken. We deemed both situations to be indicative of the Paradox.

Concerns could rise in line with the *behaviour scores*. However, if worries are already strong, they might not further increase. Furthermore, unease might be moderated by knowledge of protection. Therefore, it is unlikely that concerns will increase as much as behaviour. This should mitigate the *disparity*, hence reducing its prevalence. The prevalence is defined as the number of individuals that present the issue. Since treatment participants experience the game, only their group should be affected. For our third hypothesis, we asserted that the *disparity* will be significantly less prevalent after the gameplay session. If accepted, the Paradox might have been mitigated.

H4. We wish behaviour change to originate from our prototype, rather than our methodology. Apart from the treatment group, responses should differ little from pretest to posttest. For our final hypothesis, we posited that the control-group *disparity* will *not* be significantly less prevalent after the session. If this is accepted, there should have been little influence from our questionnaires. Most importantly, if H3 is also accepted, the game might have mitigated the disparity. This would answer our fifth subquestion and encourage the conduction of a medium-term analysis.

We formally summarise our four hypotheses below:

1. The final mean *behaviour score* of the treatment group will be significantly higher than its initial mean *behaviour score*.
2. The final mean *behaviour score* of the control group will *not* be significantly higher than its initial mean *behaviour score*.
3. The number of treatment-group individuals expressing a *disparity* in posttest will be significantly less than the number of treatment-group individuals expressing a *disparity* in pretest.

4. The number of control-group individuals expressing a *disparity* in posttest will *not* be significantly less than the number of control-group individuals expressing a *disparity* in pretest.

If the prevalence of the Paradox is reduced, this suggests short-term success. We can then use our findings to inform a native implementation. We should also receive qualitative feedback from our smartwatch owners. Based on their views, the game can be further refined. This should result in an application with longer-term feasibility.

Privacy feature selection

Ecosystem selection criteria. To evaluate behaviour, we explore the usage of privacy-protective features. Therefore, we must first identify which settings are widely available. However, this is challenged by the heterogeneity of smartwatches. Models are diverse, ranging from Fitbits to Pebbles to Apple Watches. While some might possess permissions, others lack third-party apps. Therefore, for a consistent analysis of behaviour, we needed to select a single ecosystem.

Selection was undertaken through three criteria: *functionality*, *autonomy* and *amenability*. Since we were choosing an ecosystem rather than features, we did not use the approach from Section 3.5.

1. *Functionality*: Firstly, products should possess privacy-protective functionality. Without such features, participants will be unable to protect their data. Furthermore, as we seek to explore behaviour, we would have no actions to analyse. These settings should be included in most models as default. This would provide owners with a similar opportunity for protection.
2. *Autonomy*: Secondly, devices should have some autonomy from a paired smartphone. Data is often synched between smartwatches and mobile phones. This enables call functionality and convenient notifications. However, the focus of our analysis is the smartwatch. Therefore, protective features should be native to that product.
3. *Amenability*: Finally, devices should be amenable to app development. This prototype seeks to inform a native smartwatch implementation. Therefore, for feedback to remain relevant, we should target the same ecosystem.

Ecosystem selection. Based on the above criteria, we selected (Android) Wear OS as our smartwatch environment. This system is hosted on a wide range of wearable devices. Therefore, its popularity should encourage a large sample size.

In terms of *functionality*, Watch OS offers a broad range of protective features. These include screen locks and application permissions. Through their use, individuals are able to guard their data. These features are included as default on version 1.1². Since all models can upgrade to this version, the settings are available to our participants. Therefore, behaviour can be compared fairly.

These devices also have a strong degree of *autonomy*. Whereas the Apple Watch is configured through a phone, Wear OS provides native settings. Furthermore, they offer an on-device store for downloading apps. These apps are then installed solely on the watch, without a paired program being required. This again differs from the Apple ecosystem. By exploring native behaviour, we directly analyse smartwatches.

Finally, Wear OS appeared *amenable* to app development. Third-party tools can be built and deployed directly to the watch. These applications need not be hosted on an app store, but can be ‘side-loaded’ through a USB connection. This is supported through the Android Studio IDE, a usable and full-featured environment. Apple Watches can also host third-party applications. However, they are less able to monitor the device’s settings. The deployment process is also complex and expensive. Based on these considerations, Wear OS appeared preferable.

Feature selection. Now that the ecosystem has been selected, we chose the features to analyse. Since we scoped our focus, we do not study the ‘general’ features from the comparative interviews (Chapter 6). To assess this wearable environment, all tools were specific to smartwatches. Features were selected through the consistent criteria from Section 3.5: *simplicity*, *utility*, *applicability* and *concern correspondence*. The contextualised questions are outlined within the Main Survey Design subsection.

Selected features. Based on the above criteria, we selected three privacy-protective actions. Their frequency will be analysed to evaluate participant behaviour. The features comprised: enabling a *screen lock*, disabling *GPS* and restricting *app permissions*. Our rationale is outlined below.

Screen locks are provided through three options: PINs, passwords and graphical patterns. While the feature might pose a slight delay to users, it limits unauthorised access. Passwords are found on many modern technologies. Since the concept should be understood, they comply with the simplicity criteria. They also have clear utility,

²We make reference to the Android Wear version numbers. They have recently been reset following the rebranding to Wear OS.

since they limit access to data. As smartwatches can store a range of personal details [118], unauthorised access could cause damage. If a watch is lost or unattended, passwords offer good protection.

Screen locks are also supported on all Wear OS models. This signifies their broad applicability. Finally, the feature has direct correspondence with our first contextualised question. The incident concerns a watch being accessed by a stranger. This issue belongs in our *smartwatch* threat model from Section 3.4. As with many consumer devices, watches face a reasonable risk of loss. Furthermore, as highlighted by Ricci et al. [320], their “*size and portability makes them easy to steal*”. Screen locks reduce the likelihood of access by limiting unauthorised usage. Such passwords have been suggested to reduce theft risk [102]. Therefore, if a person is concerned about the incident, a screen lock might offer protection.

GPS is beneficial for navigation and sporting activities. This service is commonly used by apps to offer personalised functionality. While the advantages of GPS are clear, it can also pose privacy risks. When the service is enabled, positions can be identified. For example, if one spends the night at a location, a company might infer where the person lives [180]. Although positions might be tracked through other networks, GPS has the greatest precision [398]. Smartwatch recordings might be even more accurate, since the device is attached to the body. We do not call for the abandonment of GPS, since it provides great functionality. However, it could be disabled when not required.

Disabling this feature matches all of our criteria. Firstly, GPS is popular and commonly used on mobile devices. Furthermore, the service is accessible through the home screen. Therefore, it should be simple to use. Secondly, there is utility in disabling this service. GPS allows geographic tracking, supporting the inference of sensitive information [180]. If one wishes to reduce their risk, disabling the service is wise. Thirdly, GPS is supported on all Wear OS smartwatches. These devices also allow the feature to be toggled. This enables an evaluation of protective behaviour. Finally, we had a direct correspondence between GPS and our second contextualised question. The incident involved a company tracking an individual’s location. When the service remains enabled, this situation is feasible [180]. Indeed, this has been demonstrated in the recent Strava controversy [174]. By disabling GPS, the risk is reduced [314]. If a user is concerned about this incident, they can limit its likelihood.

App permissions allow application data access to be limited. Smartwatches can contain an array of sensitive information [118]. Apps interact with these details to provide useful functionality. However, since users rarely read their policies [269],

they might be unaware of this access. If users wish to restrict data collection, app permissions offer a rare opportunity. We accept that these settings are ‘all-or-nothing’, in that users cannot negotiate their own conditions [35]. The menus also tend to be opaque and can impede convenient functionality [36]. However, they do prevent the collection of personal data. Such collected details are frequently exchanged with other companies [334]. Therefore, permissions could also reduce this sharing risk.

These settings met our four criteria. Firstly, permissions should be known to users, since they are included on smartphones and social networking sites. Although they can lack usability, smartwatch owners should be aware of the concept. Secondly, permissions have clear utility through protecting data. By restricting settings, the access to details can be limited. Thirdly, permissions are available on all Wear OS devices. They should therefore have broad applicability. Finally, we established a correspondence between permissions and our final contextualised question. In this incident, an app company shares a user’s information. Again, such sharing is a common occurrence within the digital economy [334]. However, if permissions are restricted, the disclosing party is limited in their access. Therefore, if a person opposes this issue, settings might reduce their risk.

Experimental structure

Overview. We wished to evaluate whether our prototype could mitigate Paradox prevalence. Therefore, we selected a pretest-posttest structure. This allowed us to analyse the issue before and after our gameplay session. An illustration of the structure can be found below in Figure 7.1.

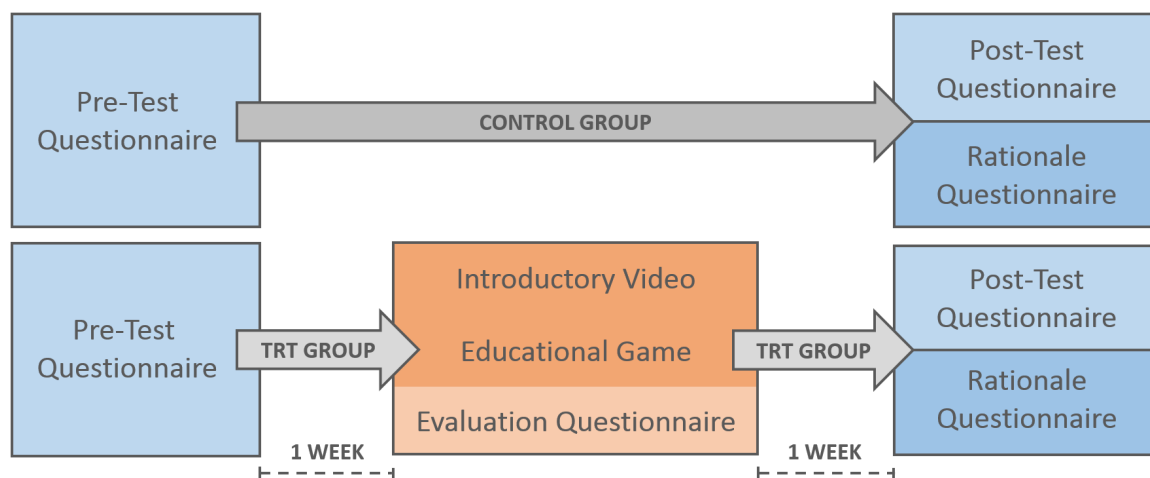


Figure 7.1: Experimental Structure

Firstly, we solicited concerns and behaviour through the pretest questionnaire. Then, one week later, our participants were split into a treatment group and a control group. While the former took part in an interactive session, the latter did not. The session was composed of activities: an introductory video, the prototype game and an evaluation questionnaire. The video highlighted privacy and demonstrated gameplay. This sought to enhance the usability of the prototype. Once this content was watched, the game was played several times.

The session was concluded by an qualitative questionnaire. Through this, we solicited detailed evaluations of the game. Based on smartwatch user feedback, we could inform our refined application. A one-week gap was then installed before the posttest phase. This period was included for two reasons. Firstly, it aimed to reduce priming, since privacy was disguised in the questionnaires. Secondly, the gaps granted individuals the time to practice new techniques. Since such periods are included in comparable studies [19, 20, 115, 223], we adopted this approach.

Finally, all participants completed the posttest questionnaire. The first half was identical to the pretest form, enabling a fair comparison. The second half differed by requesting qualitative responses. These replies were then used to dissect decision-making rationale. It was crucial that these questions did not influence the first half. Therefore, page navigation was disabled.

Methodology selection. We chose to solicit concerns and behaviour through an online questionnaire. This approach was selected for several reasons. Firstly, by hosting the questionnaire online, we could target a distinct sample. The Prolific service, to be outlined, provided access to large numbers of UK smartwatch owners. If the Paradox is prevalent in this population, we corroborate our findings from earlier chapters. Since many top-tier privacy studies use online simulations [206, 315, 383], it was deemed appropriate for a prototype. Secondly, this online approach reduced our time investment. If the prototype was tested offline, the process might have taken several weeks. It would have also been challenging to recruit such a specialised sample. We decided this time was better-invested in the final study. Finally, online prototypes provide individuals with a safe environment. Participants might have never explored their settings before. Therefore, they might be reluctant to adjust their configurations. Fortunately, our prototype supports experimentation in a risk-free context.

Main survey design

Design process. To solicit responses in pretest and posttest, we distributed a questionnaire. It was important that privacy was not primed, as this could inflate concerns

[65]. Therefore, we included a large number of decoy questions. As these queries concerned other matters, the focus of our study should be disguised. This questionnaire can be found below in Table 7.1. Decoy queries are denoted by ‘-’, while the numbers indicate the concern-behaviour correspondence. The rationale questions were only asked in posttest, and were included to extract qualitative opinions.

Once we felt that the document was complete, it received face validation [96]. As before, this consisted of a review by a privacy researcher. A questionnaire expert also advised on the mitigation of response biases. After the form was refined, it was ready for distribution. The questionnaires consisted of five sections: *Demographics*, *Concerns*, *Behaviour*, *Rationale* *Informed consent*. The final section outlined our ethical procedures, and is omitted for brevity. Since the Rationale questions were only asked in posttest, they are described in the next subsection. Our other queries are outlined below.

Demographics. As in our previous studies, we collected gender, age range and highest education level. This was undertaken so we could monitor the representativeness of our sample. We also undertook group comparisons, though these are omitted for the sake of brevity.

Privacy concerns. We then solicited participants’ privacy concerns. To mitigate the influence of order effects, these questions were placed before the Behaviour section. This was undertaken since the topic might be primed by reporting actions [65]. As mentioned earlier, concerns were gauged through contextualised questions. We followed a similar approach to our interviews, with queries grounded in a defined environment. However, since we now scoped to smartwatches, we did not reuse Chapter 6 incidents.

The incidents were drawn from the potential consequences of privacy inaction. We used the consistent selection criteria, as justified in Section 3.5. They consisted of: *relevance*, *feasibility* and *behaviour correspondence*. As previously mentioned, we consider the *smartwatch* threat model from Section 3.4. Based on the criteria, we selected three distinct issues. They were: *stranger access*, *location tracking* and *app sharing*. Their selection is justified below.

Stranger access could occur if a watch was lost or left unattended. Since devices can possess sensitive data [118], they could reveal personal details. If communication apps are installed, an owner’s identity could also be impersonated. This incident is an archetypal violation, classified as intrusion within Solove’s taxonomy [352]. Since unauthorised access is often invasive, the incident fulfils the relevance criterion.

Table 7.1: Pretest/Posttest Questionnaire
The ‘Qualitative Rationale’ Section was only asked in Posttest

#	Demographics
1	With which gender do you most identify? <i>Male</i> <i>Female</i> <i>Other</i>
2	What is your age? <i>18-25</i> <i>26-35</i> <i>36-45</i> <i>46-55</i> <i>56-65</i> <i>66 and over</i>
3	What is your highest level of education? <i>School/GCSE</i> <i>A-Level/College</i> <i>Degree</i> <i>Masters</i> <i>PhD</i>
#	Contextualised concern questions (on five-point scale from Very Pleased to Very Displeased)
-	How would you feel if your watch’s screen became dirty?
1	How would you feel if a thief viewed all the data on your watch?
-	How would you feel if your watch’s battery died during the day?
2	How would you feel if other companies were spying on your watch’s location?
-	How would you feel if your watch’s apps slowed down?
3	How would you feel if your watch shared your text messages with another company?
#	Behaviour questions (on five-point scale from Always to Never)
-	How often do you install your watch’s updates?
3	How often do you disable/refuse the permissions of watch apps?
-	How often do you clean your watch?
2	How often do you use location services (GPS) on your watch?
-	How often do you charge your watch overnight?
1	How often do you use a passcode/pattern lock on your watch?
#	Qualitative rationale
1	In which way do you believe you learn most effectively? <i>Visually</i> <i>Aurally</i> <i>Verbally</i> <i>Physically</i> <i>Unsure</i>
2	Do you think you understand how to protect your smartwatch data? Why? <i>Yes</i> <i>No</i> <i>Unsure</i>
3	Do you feel confident that you can protect your smartwatch data? Why? <i>Yes</i> <i>No</i> <i>Unsure</i>
4	Do you think taking action to protect your smartwatch data is worth the effort? Why? <i>Yes</i> <i>No</i> <i>Unsure</i>
5	Do you think your smartwatch data faces a realistic threat? Why? <i>Yes</i> <i>No</i> <i>Unsure</i>
6	What would influence you to care more about your smartwatch data?
7	What would influence you to take greater action to protect your smartwatch data?
8	Indicate your agreement or disagreement with the following statement: “If you’ve got nothing to hide, you’ve got nothing to fear”. <i>Strongly Agree</i> <i>Agree</i> <i>Neutral</i> <i>Disagree</i> <i>Strongly Disagree</i>

This scenario was also deemed to be feasible. Smartwatches might be left unattended during sports or other activities. Furthermore, they are small and therefore simple to misplace. Indeed, their “*size and portability makes them easy to steal*” [320]. Since wearables are both popular and expensive, these products may face an

additional theft risk. Finally, the incident had direct correspondence with protective behaviour. If an individual fears unauthorised access, *screen locks* offer a simple solution. Passwords have even been suggested as a threat deterrent [102]. Therefore, users can respond to their concerns with this protective feature.

Location tracking can be undertaken through GPS, Wi-Fi or mobile networks. However, the former offers the most accurate monitoring [398]. This issue has clear relevance to privacy, as it concerns the surveillance violation [352]. Companies might track individuals to target advertising or infer additional details. This data could also be shared with commercial partners, as is common in the digital economy [334].

The incident fulfils the feasibility criteria, since location is requested by a range of apps. Furthermore, GPS is enabled by default on most Wear OS devices. Since individuals tend not to adjust their smartwatch settings [371], their location might be accessible. The recent Strava tracking has highlighted this issue [174]. Finally, the issue can be mitigated through disabling GPS. This action reduces the risk of location tracking [314]. In this manner, we establish a direct correspondence between concern and behaviour.

App sharing can occur after an application accesses personal data. The access is generally undertaken to offer convenient functionality. However, apps can request unnecessary information, as found in Android ‘grayware’ [151]. Moreover, details can be synched with servers for further processing. Therefore, companies have opportunities to read sensitive data. This question has relevance to privacy, particularly the disclosure and secondary use violations [352]. Furthermore, such sharing has feasibility, as demonstrated in the recent Facebook controversy [156]. Finally, settings exist to mitigate the risk. To share data, a company must first acquire the information. When permissions are restricted, this access is at least partially limited.

Concern scores. Participants indicated their reaction to each of the contextualised questions. Responses were made on a five-point Likert Scale, ranging from ‘Very Pleased’ to ‘Very Displeased’ via ‘Neutral’. Again, we avoided the term ‘concern’ to help disguise the purpose of the questionnaire. In our draft form, the scale began at the ‘Neutral’ point. However, we feared that this might encourage participants to express concern. Based on our amended scale, we received metrics between 1/5 (low) and 5/5 (high). By taking a mean across the incidents, we calculated the *concern score*. By analysing these scores, we explored whether concerns changed as the study progressed.

Privacy behaviour. To analyse protective behaviour, we explored how often the selected features were used. They consisted of: enabling a *screen lock*, disabling *GPS*

and restricting *app permissions*. These tools should respond to participant concerns, supporting a correspondence between our questions. Again, this sought to establish the *principle of compatibility* [12]. Protective features were the main focus of our prototype game. By teaching participants how to use settings, we should encourage greater protection.

Behaviour scores. Responses were made on a five-point Likert Scale. This ranged from ‘Always’ to ‘Never’ via ‘Occasionally’. We did not include definitions for these frequencies, as we explored the subjective opinions of our sample. Based on the responses, we received metrics from 1/5 (low) to 5/5 (high). Since we analysed the avoidance of GPS rather than its use, their replies were reverse-coded. Finally, taking a mean across the three features, we calculated the *behaviour score*. By comparing scores before and after gameplay, we explored whether actions changed.

Bias mitigation. As previously mentioned, concerns can be inflated when privacy is salient [65]. Furthermore, we did not want behaviour to be exaggerated through social desirability bias [138]. Therefore, our research was not advertised as relating to privacy. Instead, the questionnaires ostensibly concerned smartwatch opinions. Since the cautious might avoid privacy studies, our approach also sought to reduce non-response bias [164].

As highlighted earlier, we included several decoy questions. They were also used in an attempt to disguise privacy. These queries concerned the general functionality of smartwatches. Participants faced another three incidents and features. They were shuffled in with the aforementioned questions. To conceal our purpose, non-privacy actions corresponded with the non-privacy scenarios. This allowed our queries to follow a consistent format.

Posttest rationale design

In the first stage (pretest) and final stage (posttest), participants completed identical questions. However, an additional section was included in the latter case. This solicited qualitative justifications for behaviour. Through this approach, we analysed decision-making rationale. This informed the development of our refined application. For brevity, we only discuss those questions with findings of relevance.

PMT factors. Our interviews, outlined in Chapter 6, highlighted many justifications for protection avoidance. Most common among these factors was a lack of awareness. We hoped that after gameplay, a participant’s risk perception would increase. This might result in privacy-protective behaviour. However, individuals

consider many matters before undertaking an action. As has been previously described, this can be modelled through Protection Motivation Theory (PMT) [328]. To evaluate behavioural rationale, we explored the matter through four questions.

The first asked participants whether they believed they understood protection. This assessed their confidence in their own abilities, exploring self-efficacy [43]. Since the treatment group had played the prototype game, we would expect them to be more confident. In the second question, we asked individuals whether they felt they could protect their watches. This also assessed self-efficacy, in addition to response efficacy. Again, those playing the game should become better-equipped to protect themselves. We then asked participants whether they considered privacy worth the effort. If not, our prototype should do more to highlight its advantages. This explored the role of perceived threats and response costs. Finally, users were asked whether they felt at risk. Through this question, we directly assessed their threat appraisal. We would expect the treatment group to gain greater awareness through the game. If not, the refined app should highlight risk exposure.

Further influence. Our game sought to provide both education and encouragement. However, we knew our prototype could be further extended. Therefore, we asked participants what would influence their behaviour. Respondents might not have perfect awareness of their own rationale. Nevertheless, they could still offer helpful suggestions. If certain factors predominated, these could be adopted in the final application.

Evaluation questionnaire design

Overview. At the end of the gameplay session, participants completed an evaluation questionnaire. This sought to gauge opinions of the interaction. Based on feedback from smartwatch owners, we could then refine the final application.

The questionnaire consisted of four sections: *Codes*, *Evaluations*, *Performance* and *Opinions*. Their questions are now outlined below.

Codes. Before completing the questionnaire, participants were required to watch an introductory video and play the game. However, compensation was allocated at the end of the form. Therefore, there was a risk that the first two components would be skipped. To test whether individuals had participated, we requested two visual codes. These codes were embedded within the video and the game. Incorrect responses were filtered from the study, ensuring findings were built on informed replies.

Evaluations. We also wished to gauge the views of our participants. In order to evaluate our game, the prototype was rated across four factors. These comprised:

enjoyment, usability, game education and session education. Firstly, respondents indicated whether they enjoyed the game. Since enjoyment contributes to immersion [305] and retention [28], we wished our prototype to be appreciated. Secondly, we solicited the usability of the application. For behaviour to be influenced, the game must be usable [209]. Thirdly, since we wished to impart knowledge, we asked whether the game was deemed educational. If not, then additional information might be needed. Finally, we analysed whether the session was educational as a whole. This also evaluated the video, which sought to introduce our game. If both components are helpful, protective behaviour might be encouraged.

Performance. Gameplay progress was analysed through two metrics: *challenges completed* and *total points*. The prototype included three challenges, with each concerning a protective feature. If participants master these tools, their protective behaviour may be supported. Points were assigned based on gameplay performance. This allowed us to monitor the difficulty of the prototype. If the game is too simple, individuals might not be immersed through ‘pleasurable frustration’ [152]. This occurs when tasks are challenging but fun. By recording these metrics, we can refine the prototype accordingly.

Opinions. Since our participants owned smartwatches, we expected them to possess relevant opinions. Therefore, we wished to refine our design based on their views. To achieve this, we asked users what they liked and disliked. Since they might possess helpful ideas, we also solicited suggestions for refinement. Finally, as in posttest, we asked what might change their behaviour. If certain factors predominate, they could be considered for the final app.

Participant recruitment

Platform selection. To refine our prototype, we required a large quantity of informed feedback. However, smartwatch owners might be challenging to recruit through traditional techniques. With these devices still being novel, public advertisements might return few responses. Therefore, we made use of a crowdsourcing platform. These services provide access to large numbers of specialised individuals. Such tools are popular in top-tier research [206, 339, 340], having been deemed valid for HCI studies [73]. Although user pools might be overrepresented in certain demographics [73], they were deemed beneficial for prototype refinement.

We selected the Prolific crowdsourcing platform³. Although many alternatives exist, including Amazon Mechanical Turk (MTurk), Prolific was selected for two

³<https://www.prolific.ac/>

reasons. Firstly, it is a British service, with most of its users resident in the UK⁴. Since we explore the Paradox in this nation, it was prudent to choose a UK platform. Secondly, to support reliable research, we sought trustworthy participants. While this can never be guaranteed, honesty has been found greater in Prolific’s user pool than MTurk’s [296]. Therefore, we opted for the UK service.

Distribution. To recruit participants, an advertisement was uploaded to the platform. This outlined the experiment, the time required and the compensation. As previously mentioned, this advert did not mention the topic of privacy. Therefore, cautious individuals should not have been deterred from participation.

If candidates fulfilled the screening criteria (described below), they received an invitation. When this was accepted, they were then directed to our questionnaire platform. To maximise usability and security, we opted for Bristol Online Surveys⁵. This site specialises in academic studies and was recommended by our ethics committee. Once participants completed the questionnaire, they received their compensation.

For the second stage, we split our participants into the treatment group and the control group. The former were invited to our session, which included the video, game and questionnaire. Again, they received compensation once the form was completed. In contrast, the latter group did not participate in these tasks. Finally, one week after the session, all users were invited to the final stage. If individuals did not answer this questionnaire, they were removed from the whole study. This was essential, as we required a complete dataset to enable comparisons.

Screening. Individuals were allowed to participate if they fulfilled three criteria. Firstly, they must be 18 or over and able to provide informed consent. Secondly, since only device owners could report behaviour, they had to possess a smartwatch. This also increases the chance that concern responses are informed. Finally, participants were required to have a Prolific approval rating of at least 90%. This was the highest filter available, and such levels are good practice in crowdsourced research [297]. By considering the views of reliable users, we sought to increase our validity.

Compensation. To incentivise participation, respondents were ethically compensated for their time. For the first stage questionnaire, individuals were given £1.30 for the 3-minute task. The treatment group were then compensated £4 for the 25-minute training session. Finally, all participants received £1 for the 4-minute posttest survey. Since we sought to disguise our purpose and remuneration was applied equally, we do not believe compensation biased our findings.

⁴<https://www.prolific.ac/demographics/>

⁵<https://www.onlinesurveys.ac.uk/>

7.3 Gameplay Session Design

General session

Overview. Our treatment group took part in the smartwatch session. It was composed of: the *introductory video*, the *prototype game* and the *evaluation questionnaire*. Since we have already described our survey, we now outline the first two components.

Introductory video. Educational games are a relatively new field, particularly for privacy [363]. Furthermore, it is likely that individuals have never used such an application. Therefore, to introduce the environment, we began with a brief video. The webpage can be found below in Figure 7.2.

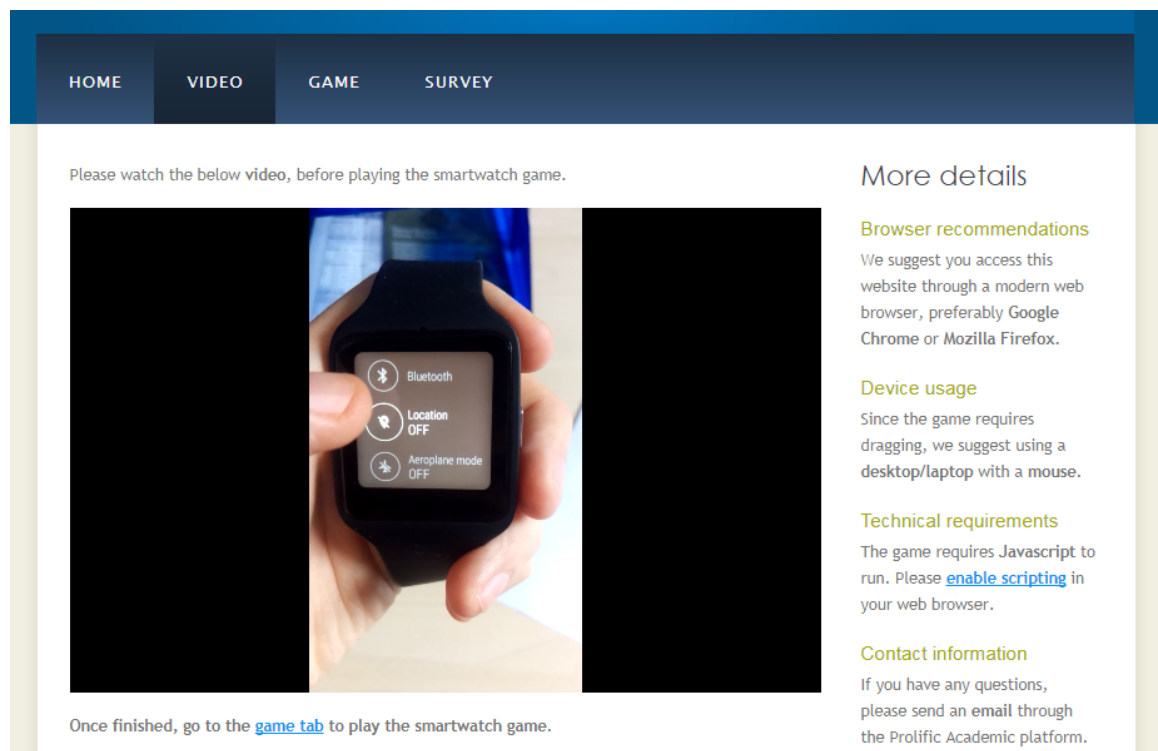


Figure 7.2: Introductory Video

The video first introduced the topic of smartwatch privacy. This placed the game in context and justified its importance. We then demonstrated how to navigate a smartwatch interface. This involved a tour of the three privacy-protective features. By learning where settings are placed, users should perform better in the game. Finally, we discussed the online prototype. In addition to describing the narrative, we provided a brief demonstration. By introducing the gameplay, we believe the app was better-placed to impart knowledge.

Prototype game

Narrative. We now outline the focus of our study: the prototype game. Two screenshots can be found below in Figure 7.3. The leftmost screen presents a privacy challenge, while the rightmost displays ordinary gameplay. For brevity, additional images can be found in Appendix B.1. This application directly informs the final game in the next chapter. Since the prototype was considered effective, many of the design approaches were retained. Therefore, to reduce repetition, Chapter 8 only highlights where the two apps differ.



Figure 7.3: Prototype Educational Game: Challenge (left), Map (right)

In our game, *(Smart) Watch Out!*, players must navigate their character to a local shop. En route, they collect coins to pay for their purchases. Their town is inhabited by two types of Non-Player Character (NPC): *villagers* and *thieves*. Villagers are friendly and ask the player for privacy advice. In contrast, the thieves trigger a timed interactive challenge. In these challenges, users must configure a simulated menu. These menus correspond with the three privacy features. This helps participants learn and practice protective behaviour. If they act too slowly, their health expires and the game is over. However, if they succeed they are awarded additional points. They then continue their journey through the rest of the map. Players win by reaching the shop, having overcome the three thieves.

This narrative was selected for several reasons. Firstly, to engage our audience, it was important that our theme was accessible. Secondly, the narrative needed to

be broadly relatable [248]. We initially considered a spy game, but this might have made privacy appear irrelevant to most users. Finally, understanding is enhanced when physical and virtual risks are aligned [150]. Through including thief characters, this alignment should be achieved.

Interaction. Navigation takes place on a top-down 2D map, similar in style to many retro games. This design was chosen for its appeal and familiarity. Players controlled the app like a smartwatch interface. They clicked to tap, dragged to swipe and clicked the side button when necessary. Context is important for gamification [321], and participants could apply lessons to their real watch.

Customisation. In educational games, customisation has been found to encourage immersion [28]. This was crucial for our prototype, as engaged users tend to retain information [28]. Therefore, players began the game by naming their character. They selected a three-digit title, similar to the style of retro games. Individuals were then invited to select an avatar, customising the character on three levels. They could toggle between options for gender, hair colour and skin colour. Avatars have also been found successful for increasing immersion [27]. We hoped these techniques would encourage participant engagement.

Questions. As previously mentioned, there were two types of Non-Player Characters. The first, villagers, ask questions related to the protective features. An example screen can be found within Appendix B.1. This approach seeks to refine the lessons from the interactive challenges. Participants are given two options with which to answer the question. If they respond correctly, they receive additional points. However, if they answer incorrectly, the correct response is highlighted. Further advice is also given, reinforcing the privacy education.

Challenges. Thieves trigger challenges, which require the user to update the simulated settings. These settings mimic a smartwatch environment and hence resemble the participants' device. Furthermore, the challenges directly correspond with the three protective features. Therefore, users are taught how to enable screen locks, disable GPS and restrict their app permissions. For example, a participant might have to navigate through the 'Settings' and 'Connectivity' menus to disable GPS. As users undertake these challenges, they might learn new behaviour.

A time limit is assigned to each privacy challenge. The quicker the task is completed, the more points are received. This encourages participants to memorise the navigation. However, if the time expires, the game is over. While this adds pressure to the experience, 'pleasurable frustration' tends to encourage immersion [26]. When confronted by a challenge, users are able to decline the task. This was important, as

enthusiasm can decrease when individuals feel controlled [38]. It also sought to simulate real life, where privacy risks can be ignored. However, if the threat is overlooked, the character's score and health are halved. Since incentives are more influential than coercion [38], we believed this to be a sensible approach.

Completion. Players complete the game by reaching the shop at the end of the map. If their score is sufficiently large, it is then added to a leaderboard. Even if the participant succeeds in all three challenges, they might not top the list. Therefore, they are encouraged to repeat the game to improve their score.

Behaviour change principles. Now that we have outlined the game, we discuss its adherence to behaviour change approaches. Since the Chapter 8 app extended this prototype, the same techniques were adopted. As previously mentioned, a lack of awareness appears to contribute to the Paradox. However, even if individuals are alerted to an issue, behaviour change is not guaranteed [38]. To support this process, awareness should be accompanied by education and practice [332]. This is another reason we selected interactive games over traditional campaigns. Finally, individuals must feel incentivised to use new techniques [119]. To encourage protective behaviour, our games addresses each of these levels.

Firstly, we increase awareness through privacy challenges. By highlighting the potential risks, individuals should be more alert to smartwatch threats. Furthermore, by showcasing protection, users should learn of the features. In the final application, a slideshow also seeks to illustrate the consequences. Secondly, education is offered through feedback and hints. Whether challenges are completed or failed, helpful information is provided. Similarly, even if questions are incorrectly answered, the correct response is explained.

Thirdly, the challenges provide an opportunity to practice new behaviour. By experimenting on a simulated interface, users become familiarised in a safe environment. Finally, protection is incentivised through scores and success. By learning new techniques, users receive points and congratulation. As participants complete tasks, they gain access to additional challenges. The final app also includes levels and difficulties to extend this experience. Privacy is portrayed as empowering, rather than an impediment to convenient functionality. Therefore, through advice and positive reinforcement, we seek to encourage protective behaviour [101].

Design principles. Since we wished our apps to be persuasive, we implemented Learning Science principles. We first subscribed to Quinn's four tenets [312]: goal-orientation, challenging, contextual and interactive. The prototype achieved this

through challenges, time constraints, rich narrative and interaction, respectively. Our refined app extended this with levels, difficulties and dynamic customisation.

The six principles of educational game design were also implemented [26]. These were: unique identity, immersion, interactivity, increased complexity, informed teaching and instruction. Our prototype accomplished this, respectively, through: avatars, rich narrative, high responsiveness, varied tasks, scores/challenges, and NPC feedback. The final app also included levels, hints and multiple difficulties. By following best practice, our games were well-informed.

Pilot study

Recruitment. While we believed our prototype would be effective, we sought additional feedback. Therefore, we conducted a brief pilot study with 30 participants. The sample size was over 10% of our final treatment group, with this adhering to recommended standards [224]. Participants were also recruited through the Prolific crowdsourcing platform. We recruited for those without wearables, since we sought distinct users from the actual study. These individuals watched the video, played the game and completed the questionnaire. If they could navigate our prototype successfully, smartwatch owners should find it usable.

Amendments. For brevity, we only discuss our resultant amendments. Firstly, the game’s visual code appeared challenging to locate. It was only found by 53.3% of participants, despite them investing sufficient time. Therefore, it was made more prominent and included in several places. Secondly, participants completed many challenges ($\bar{x} = 2.3/3$) and received high scores ($\bar{x} = 801.7/1000$). While this was promising, we sought to trigger ‘pleasurable frustration’ [152]. In seeking to encourage immersion, we reduced the time limits for the challenges. Finally, the most-frequent criticisms concerned interface usability. Participants had difficulty playing the game when using touchscreen devices. When swiping their display, the prototype would often move across the page. Therefore, we disabled scrolling and implemented on-screen keypads. We believe these refinements should improve the quality of the prototype.

7.4 Baseline Findings

Participants

Responses. 601 individuals responded to our pretest questionnaire. However, since multiple stages were required, there was a degree of attrition. We received a final

sample of 504. This contributed to an attrition rate of 16.1%. Since rates below 20% are adequate for intervention research [243], we believe our findings have validity.

Demographics. In terms of demographics, we had a varied gender mix. 55.8% were male, 43.7% were female and 0.6% identified as non-binary. This corresponds closely with the 57%-43% split from both our online survey and our interviews. Smartwatch owners have been disproportionately male [282], and this trend appears to be continuing⁶. Therefore, we should have decent alignment with our target population.

Our sample was also relatively youthful, with an estimated mean age of 31.7. 46.4% were between the ages of 26 and 35, and 26.2% were in the 18-25 group. Smartwatch owners tend to be young⁶, further supporting our sample. When considering education level, 45.6% had achieved an undergraduate degree. A further 17.2% had a Master's qualification, while 20.2% finished after their A-levels.

Duration. In terms of first stage completion, participants took a mean average of 3.5 minutes. The questionnaire was designed for three minutes and compensated as such. Since users invested additional time, it suggests their replies were considered.

Qualitative analysis

Coding frame creation. Through our study, we received large amounts of quantitative and qualitative data. Quantitative data was analysed through the techniques detailed in Section 3.2. The qualitative responses were evaluated through another robust process. We undertook inductive analysis, in the style of Ritchie et al. [323]. This approach is described in detail in Section 3.3. We produced structured indices which then evolved into our coding frames. The frames relevant to our discussed findings can be found in Appendix A.3.

Distinct frames were constructed for each qualitative question. For example, separate tables were created for the 'Liked most' and 'Liked least' queries. This supported analyses from multiple angles. We coded each qualitative query in the evaluation form and the posttest questionnaires. This resulted in a total of nine frames.

Coding approach. When a participant answered a query, they could mention several topics [307]. However, since responses were given to an online form, they tended to be brief. Therefore, most answers applied to a single theme. As before, we studied the 'comment' proportions. This format is also outlined in Section 3.3. By comparing these proportions, we could explore which perceptions are most prevalent. Furthermore, we could judge how views differ between our groups. If treatment users respond in a common manner, their views might be informed by the game.

⁶www.statista.com/statistics/739398/us-wearable-penetration-by-age/

Privacy concerns

Contextualised questions. We begin by considering privacy concerns, the first component of the Paradox. As participants had not played the game at this stage, these opinions should reflect baseline perceptions.

Stranger access was first considered, assessing reactions to unauthorised intrusion. 90.1% of respondents were (at least) displeased, with 67.3% expressing strong opposition. As only 7.3% were neutral, smartwatch owners appeared to reject invasion. These reactions could be compared against protective behaviour. Since passwords can limit unauthorised access, their usage could reduce this risk.

Location tracking was the second incident to be analysed. Although GPS is a common service, most participants opposed this issue. 91.7% expressed their displeasure, compared to only 6.2% with neutral replies. Since 68.3% were in strong opposition, this suggests many dislike geographical monitoring. This aligns with the 90% expressing concern in the Chapter 6 interviews. If a person is opposed to this issue, they can disable their GPS [314].

App sharing was our final incident. This solicited responses to the watch sharing personal details. Participants opposed this practice even more strongly, with 96.0% noting their displeasure. Although data exchanges are a modern reality, only 2.0% gave neutral replies. With 82.1% being in strong opposition, it appears the issue is unpopular. If users wish to minimise collected data, they could restrict permissions.

Concern scores. As previously mentioned, we aggregated responses to the three contextualised questions. This comprised the *concern score*, ranging from 1/5 (low) to 5/5 (high). When taking a sample-wide mean, respondents scored a strong 4.61/5. This implies that smartwatch owners claim concern for their privacy.

Privacy behaviour

Features. We now explore the frequency of privacy-protective actions. Since 90.1% opposed data access, we expected *screen locks* to be commonly used. However, only 57.2% of owners used this feature often (or more). This might be due to perceived inconvenience, since passwords can delay interactions. Alternatively, users might simply lack awareness of the feature. Our game both highlighted locks and provided opportunities for practice. Therefore, it might have the potential to encourage protection.

GPS usage was the second action to be analysed. We found that only 22.6% used this service rarely (or never). This comes in contrast to the 91.7% who opposed location tracking. GPS might be enabled for sensible reasons, such as providing

navigation. However, it might be forgotten and left on. By highlighting the potential risks, we hope individuals take greater action.

App permissions can limit the collection of personal data. Therefore, if a person is concerned about sharing, these settings could be of benefit. 96.0% disliked the final incident, with 82.1% expressing strong objections. However, only 23.8% disabled their permissions often. This inaction might be due to several factors. Configuration can be daunting for non-expert users [197]. Furthermore, as settings are often opaque [35], they might be misunderstood. Since our prototype seeks to raise awareness support practice, it might encourage protective behaviour.

Behaviour scores. We then aggregated the frequencies for all three actions. This produced the *behaviour score*, ranging from 1/5 (low) to 5/5 (high). When taking an average across our participants, the mean was 2.94/5. This suggests that far more can be done to protect privacy. With the *concern scores* appearing greater than the behaviour metrics, this might imply that the Paradox is present.

The Privacy Paradox

To analyse the disparity, we must compare each participant's responses. This is undertaken by considering the *concern scores* and the *behaviour scores*. If a person expresses strong concerns but uses protection rarely, the Paradox might be present.

Incident level. We first considered *stranger access* and the usage of *screen locks*. While most participants disliked the incident, passwords were rarely used. This led to the behaviour mean ($\bar{x} = 3.37$) being significantly less than the concern mean ($\bar{x} = 4.53$) ($Z = -12.43$, $p < 0.001$, $d = 0.85$). Although users may wish to guard their data, they might perceive protection as inconvenient.

We then analysed *GPS usage* and *location tracking*. In this case, most individuals were opposed to monitoring. However, GPS was disabled only rarely. Again, this meant that the behaviour mean ($\bar{x} = 2.64$) was significantly less than the concern mean ($\bar{x} = 4.56$) ($Z = -17.39$, $p < 0.001$, $d = 1.31$). The 'very large' effect size [333] emphasises the contrast between claims and action. This result is supported by Paradox research conducted by Menfors and Fernstedt [263]. They similarly found that while users fear tracking, they enjoy location services. In such cases, individuals might appreciate the benefits without recognising the risks.

Finally, we considered *app sharing* and *app permissions*. As before, participants expressed strong concerns. However, settings were adjusted infrequently. This might be due to the abstract nature of these settings [335]. Accordingly, the behaviour mean ($\bar{x} = 2.82$) was again significantly less than the concern mean ($\bar{x} = 4.75$) ($Z = -18.10$,

$p < 0.001$, $d = 1.39$). With this effect size also being ‘very large’ [333], the Paradox might be prevalent. Permissions may be loosened through a short-term necessity. If they are then forgotten, personal data might remain accessible.

Aggregate level. We then analysed the aggregate metrics, comparing the pretest scores. As expected, the *behaviour scores* were significantly less than the *concern scores* ($Z = -18.60$, $p < 0.001$, $d = 1.45$). This implies that while respondents oppose violations, they often do little to reduce their likelihood. Such a result corroborates the disparities discovered in Chapters 4 and 6. A heatmap displaying the distribution of responses can be found in Figure 7.4.

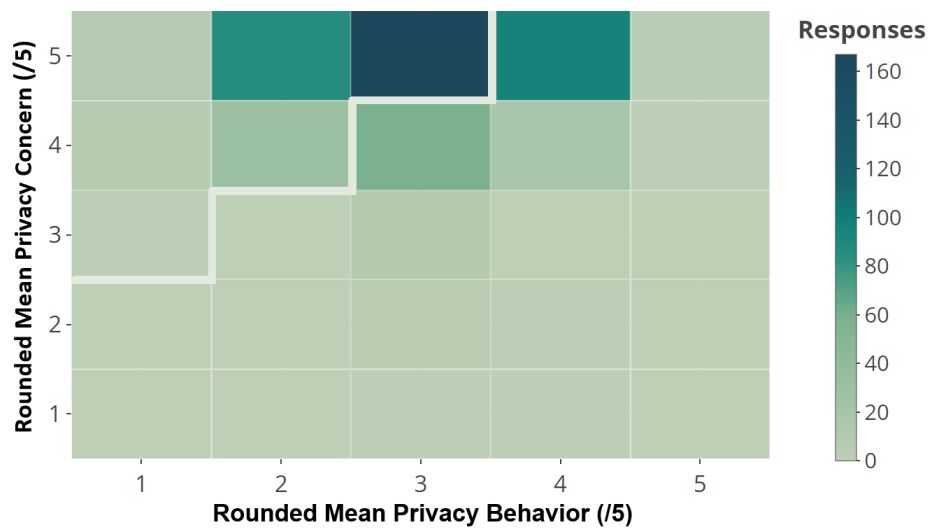


Figure 7.4: Heatmap: Concern Scores vs Behaviour Scores

To produce the graphic, aggregate scores were rounded to the nearest integer. With the *disparity* defined as a two-point gap, the top-left region denotes this issue. As shown, a large number of participants expressed their concern. However, most used protection on a sporadic basis. If this is due to a lack of awareness, hopefully the prototype will encourage protective behaviour.

Participant level. We finally studied the *disparity* at an individual level. With concerns and behaviour particular to each user, this assessed the prevalence of the Paradox. Based on the 2-point definition, 43.3% (218/504) of our participants displayed the issue. While we hesitate from comparing samples, this appears greater than the 30% found in IoT interviews (Chapter 6). This might support our assertion that wearables are particularly prone. Our prototype seeks to mitigate this prevalence.

7.5 Gameplay Session Findings

Participants

Responses. We now move on to consider our results from the second stage. At this point, a randomised half of our sample (the treatment group) took part in the prototype session. This randomisation was not stratified by demographics, but we believed this acceptable for such a large sample. Because of this division, 252 individuals trialled the game. Excluding two outliers, the mean stage duration was 35.00 minutes. This was greatly in excess of the compensated 25 minutes. Therefore, users should have invested sufficient effort.

Randomisation. We explored how pretest concerns and behaviour differed between the groups. If randomisation was insufficient, differences might owe to factors other than the prototype. Fortunately, there was no significant differences in terms of *concern scores* ($p = 0.070$) or *behaviour scores* ($p = 0.325$). Disparity prevalence also failed to differ ($p = 0.281$), as did the size of the disparity ($p = 0.715$). Based on these results, we concluded that our groups were sufficiently randomised.

Evaluation questionnaire

Evaluations. As previously mentioned, participants watched a video, played the prototype and then completed a questionnaire. This form was used to both rate the game and provide feedback. We begin by discussing the *Evaluation* section, which requested ratings across four factors. These consisted of: enjoyment, usability, game education and session education.

We first asked participants whether they enjoyed the game. This pleasure is important, since intrinsic motivation is more influential than extrinsic rewards [331]. 65.4% agreed with the statement, compared to only 60% in our pilot study. The difference was not significant, but comparisons were impeded by the disparities in sample size ($n = 252$ vs $n = 30$).

Usability was then explored, with this crucial for an immersive experience [26]. 72.2% agreed with the statement, compared to only 60% in the trial. This is encouraging, particularly since several issues were reported in the pilot study. The prototype was also found educational (79.0%), as was the session as a whole (85.7%). All distributions differed significantly from the neutral median (all $p < 0.001$). This suggests that the experience was well-received.

Performance. We now consider the gameplay performance. Our treatment group completed a mean average of 2.1/3 challenges. This appears less than the 2.3

found in the pilot study. When considering points, participants scored a mean of 648.5/1000. While this indicates good progress, it also appears lower than in the pilot ($\bar{x} = 801.7$). Increased difficulty should contribute to greater immersion [152].

Opinions. To refine the prototype, we asked participants for their qualitative opinions. In terms of what was most liked, 24.3% of comments praised the educational content. This is promising, as it suggests that users enjoyed learning about protection. Users also appreciated the amusing theme (9.8%) and the game’s simplicity (7.4%). Since these elements are popular, they will be retained in the final app. Representative quotes, along with the participant number, are provided below:

“It was a cute way to learn about privacy” (#222).

Individuals then reported what they disliked about the game. Despite refinements from the pilot study, 50.4% criticised the usability. Comments cited both the swipe-drag interaction (26.2%) and the slow speed of navigation (13.5%). We knew dragging might be burdensome, but sought to mirror a smartwatch touchscreen. Since the final app will be implemented natively, these issues should not occur.

“All the clicking and dragging wasn’t great” (#181).

When suggesting refinements, 29.7% believed the interface should be more usable. In response, we intend to implement a number of improvements. Firstly, to reduce the chance of confusion, greater feedback will be provided in challenges. Secondly, the responsiveness of the interface will be increased. Finally, since the app will be hosted on a real watch, swiping should be more natural. In terms of other suggestions, 10.2% wanted the game to be extended. Therefore, we will also introduce more challenges and questions. Since smartwatch owners inform our refinements, we believe the experience will be improved.

“Longer game with more tasks” (#266).

Finally, we asked our group what would convince them to change their behaviour. The most common factor, cited on 31.0% of occasions, was personal risk. Indeed, they would take threats seriously if their own data was in danger. To demonstrate the potential consequences, a customised app might be beneficial. Additional information was also requested (17.6%), suggesting this might be persuasive. If education is appreciated, further details can be included in the final game.

“Knowing my data and security was at risk” (#47).

7.6 Posttest Findings

One week after the gameplay session, both groups completed the posttest questionnaire. The form takes 4 minutes to complete, and was compensated accordingly. Since users took a mean of 6.13 minutes, they appeared to invest sufficient effort. Within this section, we compare pretest and posttest data in detail. For a convenient overview of the mean scores, see Table 7.2 below.

Table 7.2: Pretest-Posttest Scores

	Treatment Group		Control Group	
	Mean Concern Score	Mean Behaviour Score	Mean Concern Score	Mean Behaviour Score
Pretest	4.60	2.91	4.61	2.97
Posttest	4.57	3.15	4.60	3.03

Privacy concerns

Contextualised questions. We now consider updated reactions to the privacy incidents. Since the control group did not play the game, we expect their views to be consistent. The prototype might either provoke concerns or provide comforting knowledge.

Stranger access was the first incident, and it concerned unauthorised watch usage. Across both groups, 93.0% expressed their displeasure at the scenario. This was compared to only 3.2% who provided a neutral response. These proportions suggest that such access is still objectionable. When considering groups, the game appeared to have some influence. Concern significantly increased for treatment participants ($Z = -2.04$, $p = 0.042$, $d = 0.13$). In contrast, the control group failed to differ. Since our game highlighted the consequences, perhaps the risk became better understood.

Location tracking was the second incident to be analysed. 90.7% of our participants were in opposition, compared to the 5.8% that were neutral. No significant differences were identified for either group. Most participants appeared displeased in both the pretest and posttest questionnaires. If users wish to mitigate tracking, they can choose to disable their GPS [314].

App sharing related to an application sharing personal data. In the pretest this was most concerning, with 96.0% expressing their opposition. After the session, 93.1% reported their displeasure. Although this proportion appeared smaller, no significant

differences were found. Since users still opposed this incident, they might limit their exposure by restricting permissions.

Concern scores. To analyse views at a high level, we calculated the *concern scores*. These were aggregated from the contextualised question reactions. When studying the control group, we received a mean average of 4.60/5. This failed to differ ($p = 0.814$) from the pretest score ($\bar{x} = 4.61/5$), suggesting our methodology did not prime concerns.

The treatment group had a posttest mean of 4.57/5. This also did not vary from the pretest metric ($p = 0.544$). Such a consistency could be due to several factors. Concerns might already be maximal, therefore limiting further increases. Alternatively, if the game was educational, protective knowledge might moderate apprehension. It appears that individuals still express strong concern. However, for the Paradox to lose prevalence, protection increase in frequency.

Privacy behaviour

Features. *Screen locks* were now used often (or more) by 63.7% of our participants. This was compared to only 57.2% in the pretest stage. The frequency significantly increased in the treatment group ($Z = -3.54$, $p < 0.001$, $d = 0.22$), which might indicate that the game had influence. This feature can respond to the *stranger access* concerns. The control group were static, failing to vary from their pretest levels. This further suggests that privacy was not primed.

GPS disabling could limit the extent of location tracking [314]. Despite these benefits, the usage frequency failed to change. Only 21.0% of the treatment group used this service rarely. This did not differ from the 22.6% identified in pretest. Even after learning of the risks, users still appreciate GPS functionality. They might consider the advantages of the service to outweigh the potential issues. Fortunately, behaviour also failed to vary in our control group (21.0%). Therefore, our questionnaire should have added little bias.

App permissions seek to limit the degree of data collection. Usage was infrequent before the session, with only 23.8% using them often. Although still a minority position, this increased to 35.7% for the treatment group. This significant change ($Z = -3.65$, $p < 0.001$, $d = 0.23$) might have been encouraged by our prototype. Once again, control group behaviour failed to differ. While our results are promising, many failed to adjust their permissions. This might be because their data was deemed innocuous. Previous work [35] has found that individuals have little sense information's intrinsic value. If our refined app can explain the risks, owners might be more persuaded.

H1. We now turn our attention to our first hypothesis. In this conjecture, we asserted that the final mean *behaviour score* in the treatment group will be significantly higher than its initial score. If accepted, this might suggest that the prototype is persuasive. When comparing the pretest and posttest metrics, we did discover a significant difference ($Z = -4.40$, $p < 0.001$, $d = 0.28$). Therefore, we accepted this hypothesis. Although the effect size was small, our prototype served as a proof of concept. As our approach appears feasible, we can now refine the final application. Context has been found to be highly important in educational games [321]. Since this app will be implemented on a smartwatch, we expect it to be more persuasive.

H2. Our success cannot be analysed without consideration of the control group. The second hypothesis asserted that their final mean *behaviour score* will not be significantly higher than their initial metric. If true, our questionnaires should not have influenced actions. The control group scores failed to differ ($p = 0.150$) from pretest to posttest. Since H1 was accepted, our prototype might encourage protection.

The Privacy Paradox

An increase in protection does not equate a mitigation of the Paradox. To evaluate whether the issue has been affected, we consider both concerns and behaviour. Our comparisons are made on a per-participant basis, ensuring that these factors inhabit the same context. Before the gameplay session, 43.3% of our users displayed the issue. If this frequency is reduced, the prototype might offer some promise.

H3. We began by considering the perceptions in our treatment group. Our third hypothesis posited that their *disparity* would be significantly less prevalent after the session. Prevalence did decrease significantly ($X^2(2) = 17.58$, $p < 0.001$, $V = 0.19$), and therefore we accepted this conjecture. This may imply that our game was persuasive. However, since 29.4% of the group remained susceptible, additional work is required. By refining our prototype, the final app should have greater success.

H4. For our final hypothesis, we analysed the matter in the control group. It asserted that their *disparity* would not be significantly less prevalent after the session. If accepted, this further suggests that our questionnaires did not our findings. 39.6% now displayed the issue, contributing to a prevalence that failed to differ ($p = 0.801$). Since we also accepted H3, the prototype appears to be persuasive. Therefore, we positively addressed our fifth subquestion: *Can we mitigate the short-term prevalence of the Paradox on smartwatches?*

Comparison. The contrast between groups is illustrated in Figure 7.5. As shown, both groups began with a similar degree of *disparity* prevalence. However, after interacting with our prototype, treatment users became less susceptible. Since we studied a large quantity of smartwatch owners, we have confidence in our app’s potential.




Figure 7.5: Comparison of Privacy Paradox Prevalence

7.7 Participant Rationale

While Paradox prevalence appears reduced, behavioural rationale is currently opaque. To understand participant justifications, we analysed the second half of the posttest questionnaire. Based on this feedback, we also developed further refinements.

Through a consideration of Protection Motivation Theory (PMT), we assessed why users did (or did not) change their behaviour [328]. This analysis also allowed us to explore the efficacy of our prototype. The app sought to increase self-efficacy (through allowing practice), highlight response efficacy (through providing information), enhance risk perception (through privacy education) and reduce perceived response costs (through demonstrating simplicity). If participants are still not confident in protecting their data, we can highlight areas for refinement.

Did they understand protection?

Responses. We asked participants whether they believed they understood protection. This analysed if they had confidence in their own knowledge. Through these responses, we assessed the self-efficacy PMT component.

Our treatment group believed they were informed, with 75.4% agreeing with the statement. This was significantly greater than in the control group ($X^2(2) = 42.00$, $p < 0.001$, $V = 0.29$), where 48.0% responded similarly. The game might supported self-efficacy, one of the PMT coping mechanisms. If individuals understand protection, they should be more likely to use it.

Rationale. Participants then supplied qualitative rationale for their response. In the treatment group, our game appeared to have some influence. Of those who understood protection, 31.9% had privacy experience or knowledge. This experience might have been supplied through the interactive challenges. In contrast, the largest control-group justification was prior expertise (25.6%). Such expertise could not owe to gameplay, since this group did not participate. Therefore, their understanding was reliant on a technical background. This suggests that training approaches might be required for non-expert users. Representative quotes, along with group and participant number, are displayed below.

“I know to disable GPS and to lock to stop data falling into the wrong hands”,
(#217, Treatment).

When justifying negative responses, both groups had similar rationale. Of those who lacked understanding, 55.7% blamed this on insufficient knowledge. This indicates that existing approaches are not a panacea. As shown above, a quarter of our treatment group still lacked confidence. By extending our game, as suggested by some participants, we could include additional information.

“I’ve never kept myself informed so I don’t think I know how”, (#196, Control).

Did they feel confident?

Responses. We next asked participants whether they felt confident they could protect their watches. This further assesses self-efficacy, in addition to response efficacy.

Once again, the treatment group expressed confidence in themselves. 64.3% responded positively, compared to only 44.0% of controls. This significant difference ($X^2(2) = 21.78$, $p < 0.001$, $V = 0.21$) might suggest that the prototype encouraged assurance. If a person is confident, they should be more likely to take action.

Rationale. In terms of positive responses, justifications were similar in both groups. Confidence was most frequently expressed since the user had protected their data (29.8%). This appears rational, as features might be trusted if used often. Participants also claimed to understand protection (20.7%) and possess expertise

(9.1%). Based on this feedback, it appears that knowledge contributes to confidence. To further enhance self-efficacy, we will include additional information in our app.

“I passcode the watch and disable certain features”, (#335, Treatment).

We then considered the rationale for negative responses. In the control group, the largest justification was a lack of skill (51.7%). This appeared to emerge since: a) many believed their knowledge was insufficient, and b) protection was considered too complex. Since these individuals did not play the game, baseline awareness appears to be low. In the treatment group, some thought perfect security was infeasible (67.9%). Products were deemed vulnerable, while adversaries were considered to be too skilled. If users lack confidence in protection, they might not invest the effort. To counter this, our final app will highlight the efficacy of tools and the capabilities of attackers.

“I don’t know enough about how to protect the data”, (#64, Control).

Did they consider privacy worth it?

Responses. In our previous studies, participants often doubted the importance of privacy. For example, many interviewees (Chapter 6) justified their behaviour on short-term necessity. Others might neglect their settings if they consider them inconvenient. To explore the prevalence of these views, we asked whether privacy was worth the effort. This assessed the perceived response costs, another PMT component.

Our treatment group believed that privacy was worth the effort. 86.1% responded in the affirmative, compared to 82.1% of the controls. Proportions did not differ significantly, likely because the most were in favour. Although agreement might have been encouraged by social norms [138], response costs do not seem problematic.

This opinion might have some relationship with protective behaviour. Those who responded positively were significantly less likely to display the disparity ($X^2(2) = 6.05$, $p = 0.049$, $V = 0.11$). If individuals appreciate the advantages, they should feel encouraged to act.

Rationale. Justifications were similar across our two groups. 31.8% appreciated that protection reduces their degree of risk. If features are used frequently, data collection should be limited. Another 19.6% valued that it guarded their personal details. Based on these responses, many recognised the benefits of privacy.

“Protection better than cure, so it is well worth the effort”, (#19, Treatment).

Although most responded positively, some doubted the advantages of privacy. Of these, 57.1% did not recognise a threat. This was primarily because data was considered innocuous (45.7%). Such a result further supports the influence of risk perception (Chapter 6). If details are thought to be non-sensitive, protection might be seen as unnecessary. However, since privacy awareness is often low [99], users might not understand data processing. The refined app will take additional effort to highlight risk exposure.

“Don’t have much worth knowing”, (#210, Control).

Did they perceive a threat?

Responses. Our previous three questions explored the coping appraisal half of PMT. To assess threat appraisal, we asked participants whether they felt at risk.

Our control group were doubtful, with only 32.1% responding in the affirmative. This might explain why protective features were so rarely used. As polls suggest that only 10% understand privacy practices [99], it is unsurprising that risks are overlooked. In contrast, a threat was perceived by 49.2% of treatment participants. This significant difference ($X^2(2) = 15.27, p < 0.001, V = 0.18$) might suggest that our game highlighted risks.

Rationale. For positive responses, justifications were similar between our groups. The most popular reason, given by 20.4%, was that adversaries can launch sophisticated attacks. When users believe they are vulnerable, they might act to protect themselves. Many also regarded their data as sensitive (20.4%), which indicates a degree of awareness. By increasing awareness, we seek to highlight the risk.

“There is valuable information there which would be of interest to some people”, (#251, Treatment).

Since most of the treatment group doubted the threat, it was important to understand their rationale. Unlike many above, a large proportion doubted their data’s importance (36.3%). It appears that users either value their details or doubt their sensitivity. These evaluations might differ based on privacy awareness [350]. Another 18.9% believed that they were adequately protected. In many cases, they trusted the quality of their smartwatch vendor. Manufacturers might provide security, but data is routinely collected through apps and settings [334]. Therefore, for greater protection, we suggest the usage of privacy features.

“I don’t see why anyone would want my data”, (#59, Control).

What would encourage protective behaviour?

Based on the qualitative responses, we extracted techniques to refine our application. However, to have greater impact, we must understand what is considered persuasive. To achieve this, we asked participants what would change their behaviour.

Rationale. The most popular suggestion, given by 29.6%, was that they needed to experience personal risk. This supports the findings of both our gameplay questionnaire and the Chapter 6 interviews. By providing awareness and practice, the Paradox appears to have reduced. However, many seem reluctant to use protection unless risks are relevant. If participants have not encountered a threat, they may underestimate its probability [370]. By highlighting the risk exposure of each user, the consequences may appear more likely.

“If I felt there was a real threat to my personal data, I would take action”, (#24, Treatment).

7.8 Implications

Paradox. Our prototype game appeared to encourage privacy-protective behaviour. Those in the treatment group began to use more protection. In contrast, our control participants failed to significantly differ. Since concerns remained static, the prevalence of the Paradox decreased. We believed education and practice would be persuasive, and this might be the case. By reducing the disparity over a short term, we addressed our fifth research subquestion. As users now used greater protection, their privacy risk should have decreased.

Refinement. However, to mitigate the Paradox over a longer period, our game must be refined. These refinements are directly informed by the results of this study. In the gameplay questionnaire, we requested opinions of the prototype. It was generally regarded as enjoyable, usable and educational. In terms of appreciated elements, users liked its theme, simplicity and informative content. These concepts will be retained for the final app. We also asked users what they thought could be improved. Most highlighted the interface controls, which were impeded by the online implementation. By enhancing our app’s usability, we hope it will prove more persuasive.

To further inform our final study, we analysed actions through Protection Motivation Theory [328]. Those playing the game appeared to show greater self-efficacy and response efficacy. They also seemed more able to perceive the privacy risks. If we did influence these components, protection might have been encouraged. Based

on qualitative responses, we also analysed rationale. A range of factors appeared to influence their decisions. Some lacked privacy knowledge, while others underestimated data sensitivity. This further emphasises the influence of awareness and risk perception. To encourage protection, our refined app should address these factors. Since we doubt concerns will sharply increase, this should mitigate the Paradox.

We finally asked participants what would adjust their behaviour. Now that awareness was increased, the largest proportion cited a perception of risk. Privacy can lack salience, particularly since it is often regarded as a secondary goal [188]. Indeed, less than half of our sample claimed to feel under threat. If risk exposure is highlighted, the perceived likelihood of violation should increase [370]. This might encourage greater protection. As outlined, the findings of this study directly inform the design of our smartwatch app.

Next steps. Our final app was refined based on several considerations. Since many individuals doubt the threat, we seek to highlight their risk exposure. This might be achieved based on their current apps or behaviour. We will also include a greater quantity of information. This is key, since some users continue to feel uninformed. Such education could be delivered through customised feedback or in-built videos. As a further refinement, we wish to extend the length of the game. Participants might currently lack experience in using privacy-protective features. By including additional challenges, they have more opportunities to practice techniques. Finally, the updated app will be hosted on a smartwatch. This should mitigate many of the usability issues of the online prototype. Furthermore, since context is important for both privacy [278] and educational games [321], the app might prove more persuasive. Informed by our Chapter 7 findings, we now seek to address our central research question⁷.

⁷Central question: *Can the Privacy Paradox be mitigated in the context of smartwatches?*

Chapter 8

Can we mitigate the medium-term prevalence of the Privacy Paradox on smartwatches?

8.1 Introduction

Background

Our Chapter 7 research suggested the feasibility of smartwatch games. Through an online implementation with 504 users, we appeared to encourage protective behaviour. Since concerns remained largely stable, the prevalence of the Paradox was mitigated. This implies that our approach, reducing a lack of awareness, was appropriate. Through this research, we addressed our fifth research subquestion.

The application was constructed as an initial prototype. Since it was hosted online, we could trial the game with a large sample of smartwatch owners. The evaluation questionnaire also allowed us to collect a considerable quantity of feedback. Based on this information, we can refine this design for an enhanced app. The app will seek to encourage protective behaviour over a longer period. We believe concerns will remain stable or be moderated by additional knowledge. By assessing participants through a two-month longitudinal study, we can assess whether the Paradox is mitigated. This addresses our final research subquestion: *Can we mitigate the medium-term prevalence of the Privacy Paradox on smartwatches?*

Motivation

Improvements. While our previous findings were promising, the work could be expanded in several regards. These opportunities motivate the conduction of an empirical longitudinal study.

Firstly, the prototype was hosted as an online platform. This served its purpose of supporting a large-scale analysis. Since protective behaviour appeared encouraged, we can proceed now with a smartwatch implementation. Furthermore, the game allowed us to solicit feedback and update our design. These findings now inform our development of a smartwatch app. As context is important for educational games [321], we believe this implementation will be more persuasive.

Secondly, retention was only assessed over a one-week gap. While we considered a longer period, this length was deemed appropriate for a prototype. These one-week periods are common in security and privacy research [20,115,223,389]. However, once our game loses salience, smartwatch behaviour might revert. To explore how actions vary over time, we now opt for a longitudinal study. If individuals continue to protect their privacy, the game might be persuasive. Although previous longitudinal work has assessed privacy [203,362], this is the first to target the Paradox.

Thirdly, our control group could mitigate additional factors. In our prototype study, we explored the responses of two groups. While treatment participants played the game, the control group used no application. Since these latter individuals failed to change, we believe we limited confounding factors [344]. However, this design faced a risk from the Hawthorne Effect [7]. This describes when individuals change their behaviour simply due to interaction. For our final study, we do not wish this issue to have influence. Therefore, all participants will receive an interactive game. But whereas some will use a privacy app, others will play a generic version. If treatment protection surpasses the control group, it should not be due to the Hawthorne Effect.

Fourthly, we require empirical data to analyse behaviour. Thus far, the frequency of actions have been self-reported. This is discussed within our research critique (Section 9.3). Reporting can be prone to bias, since participants might exaggerate their behaviour. We did implement several techniques to mitigate this risk. For example, we disguised the theme, added decoy questions and provided anonymity. Furthermore, with less than 40% claiming to use permissions often, protection does not appear to be greatly exaggerated. However, for an extensive analysis of the Privacy Paradox, empirical data is required. By comparing real behaviour and concern over time, we can judge whether the disparity is mitigated.

Finally, risk awareness could be further increased. In our prototype game, we sought to educate and encourage our participants. Although we were successful, respondents identified other considerations. Many did not perceive a privacy threat, and therefore regarded protection as unnecessary. Furthermore, when asked what

was influential, most reported the experience of risk. This implies that in addition to providing education, we should also highlight exposure.

Approach. In response to the above considerations, we designed a rigorous study. We developed an (Android) Wear OS app, informed by the feedback from our prototype. To highlight risk exposure, it dynamically customised its challenges around empirical behaviour. For a robust evaluation, it was then assessed over a two-month period. This supported a medium-term Paradox exploration.

In this evaluation, we study the concerns and behaviour of 10 participants. They receive a new smartwatch and one of two applications. The treatment half receive a privacy game, while control participants get a generic version. We then monitor empirical behaviour to analyse how actions vary. Our Chapter 7 findings informed the development of this app. We continue by describing our longitudinal study. As previously highlighted, our experiment is described in this manner since it employs “*repeated measures to follow particular individuals over prolonged periods of time*” [80]. It also has similarity to other privacy longitudinal studies [384]. Alternatively, it could be termed a ‘field study’.

8.2 Methodology

Research questions

Concern metrics. We explore whether the Paradox can be mitigated over a two-month period. This was deemed to be medium-term, since we did not consider mere days or weeks. Indeed, much previous work has assessed behaviour over short periods [20, 115, 223, 389]. However, as actions were not monitored for years, it could not be classified as long-term. Since we analyse the Paradox over two months, from January 2018 to March 2018, we believe we address our final research subquestion.

To perform this analysis, we compare concerns and behaviour before and after gameplay. Furthermore, to evaluate our games, we judge our treatment group against our control group. For a fair Paradox assessment, we explicitly define our metrics.

As in Chapters 6 and 7, concerns are evaluated based on reactions to privacy incidents. Through the use of these contextualised questions, we seek to assess opinions without priming privacy. Since our focus is scoped to smartwatches, we adopt the *smartwatch* threat model (Section 3.4). Responses were made on a five-point Likert Scale, ranging from ‘Indifferent’ to ‘Very Concerned’. This was adapted from the approach of Lee et al. [232], who also studied wearable concerns. We opted for ‘Very Concerned’ rather than ‘Very Upset’, since we explicitly targeted this concept.

The incidents are described later in this section. Similar to the previous chapter, they considered *stranger access*, *location tracking* and *app access*. However, since we now evaluated a particular model (Huawei Watch 2), the scenarios were updated to reflect its capabilities. Participants responded to two questions on each topic, supporting the triangulation of opinion. We then took a mean average of each pair, appropriate since replies were aggregated over related queries [58, 78, 280]. This produced three metrics: the *stranger score*, the *location score* and the *app score*. These ranged from 1/5 (low concern) to 5/5 (high).

If a person indicates that they are concerned by an incident, they are, by definition, claiming concern. This can then be compared against their use of protective features. Concerns were solicited both before gameplay (pretest) and afterwards (posttest). By evaluating the scores in our two groups, we can analyse how opinions vary.

Behaviour metrics. Behaviour was explored through empirical data. Each smartwatch had a tracking app installed, which monitored their settings for 52 days. Watch selection is outlined later in this section, while app design is described in Section 8.3. Halfway through the study, participants received a smartwatch game. By comparing pretest and posttest behaviour, we can explore its potential impact.

As in Chapter 7, we focused on the three privacy features: enabling *screen locks*, *disabling GPS* and restricting *app permissions*. These were chosen since: a) they were previously-selected through a robust process; b) they are beneficial Wear OS settings; and c) their usage was supported by the prototype. They also corresponded with our contextualised concern questions, enabling fair comparisons. Appropriately, we developed three behaviour metrics: *lock scores*, *GPS scores* and *permission scores*.

Lock scores were generated by considering how often a password was enabled. Settings were logged every five minutes, allowing percentages to be calculated over pretest and posttest periods. Therefore, the score ranged from 0% (no usage) to 100% (constant usage). For example, if a lock was enabled in 95% of logs, the metric would be 95%. As for all features, we only received the setting status (e.g., on/off), rather than the value (e.g., the PIN). This was important for ethical monitoring. The frequency scale was considered an appropriate measure for analysis. If this feature is enabled often, there are few opportunities for unauthorised access. This reduces the risk from the first incident, *stranger access*. In contrast, if it is used rarely, there are more chances for invasion.

GPS scores were calculated in a similar manner. We considered how often the service was enabled, ranging from 0% (never) to 100% (always). For example, if it is on in 30% of our 5-minute logs, the score would be 30%. Since GPS disabling might

limit tracking, a lower metric was used to denote greater protection. A frequency scale was also appropriate in this situation. If the service is never used then monitoring is limited [314]. This reduces the risk from the *location tracking* incident. In contrast, if it is left on constantly, monitoring is a greater possibility. By assessing usage throughout the study, we can explore how behaviour varies. If this *GPS score* decreases, our game might have encouraged protection.

Permission scores were calculated through a greater context. Some permissions are innocuous, and hence pose little risk. For instance, WAKE_LOCK is required to keep the screen illuminated. Since they do not concern privacy, it would be irrelevant to analyse the acceptance of such options. Indeed, since acceptance would be near 100%, behaviour would be unfairly portrayed. However, two permissions concerned particularly private details. ACCESS_FINE_LOCATION allows an app to track the user's precise position. READ_SMS permitted an application to read a person's text messages. Access to either resource could reveal sensitive personal data. This risks aligns with our third privacy incident: *app access*.

If an app monitored precise location, the participant could be geolocated to several metres. This might enable others to determine their home address [180]. Text messages can contain sensitive details, particularly if received from friends. By permitting this access, users are placing great trust in an app. The *permission score* comprised the average acceptance percentage. For example, five of these permissions might be requested across three apps. If one is granted, the rate would be 20%. By taking a mean of the 5-minute rates, we produced pretest and posttest scores. While all permissions might be accepted at the start, some might be later constrained. If the *permission score* decreases in this manner, the game might be persuasive.

Although some apps require location (e.g., Google Maps), others might use it primarily to collect data [151]. Unfortunately, research suggests that few users recognise the difference [343]. It would be preferable to consider the risk and capabilities of each individual application. Permissions are contextual and might be used in varying ways to suit the circumstance [335]. However, with thousands available on the Google store, this was deemed impractical. Furthermore, static analysis tools might be required to identify the data flows. Therefore, the *permission score* was based on the average acceptance percentage. While this enabled behavioural comparisons, we then explored the context in posttest interviews. As expected, participants claimed to trust some applications more than others. Some reported restricting only certain apps, as discussed in Section 8.7. Through our quantitative and qualitative data, we performed a rich exploration of permission acceptance.

Collected data. In this longitudinal study, we collect a large quantity of data. The exact details will be described as this section continues. We monitor 52 days of empirical smartwatch behaviour, with settings logged every five minutes. To assess concerns, participants complete questionnaires in pretest and posttest. In addition to Likert responses, they provide qualitative comments to justify their answers.

Halfway through the study, participants interact with our smartwatch game. During this period, we monitor performance in completing the tasks. Through this, we can analyse which privacy settings are found the most challenging. After playing the games, participants also complete an evaluation questionnaire. This reviews their experience, and allows them to specify further refinements.

Finally, at the end of the posttest phase, we conduct semi-structured interviews. These 40-minute discussions were included for two reasons. We firstly evaluated the retention of privacy-related information. We also explored the rationale behind our participants' decisions.

Research questions. In this study, we seek to address the following questions.

1. *Do smartwatch users take action to protect their data?*

In previous chapters, privacy settings appear to be rarely used. However, by empirically monitoring behaviour for two months, we can explore whether this is true. If so, this has implications for privacy risk and smartwatch interface design. The results could corroborate the findings of our preceding chapters.

2. *In smartwatch environments, does the Paradox appear to be present?*

Our interviews (Chapter 6) suggested there was a frequent disparity between concern and behaviour. This was supported in our prototype study (Chapter 7), where participants rarely protected their watch. Through empirical data, we can accurately monitor behaviour. If actions are not commensurate with concerns, the Paradox may be present.

3. *Can the smartwatch game encourage privacy-protective behaviour?*

Since the online prototype appeared persuasive, we expect the native app to have success. If protection increases in the treatment group, perhaps our game provided encouragement. The app might then serve as an interactive complement to awareness campaigns.

4. *Can the smartwatch game mitigate the Paradox over the medium term?*

This query directly corresponds with both the sixth subquestion¹ and our central research question. In Chapter 7, Paradox prevalence appeared to decrease after gameplay interaction. To assess this over a longer period, we now analyse concerns and behaviour in pretest and posttest. If the Paradox can be successfully mitigated, this offers promise to privacy applications.

5. *What factors influence smartwatch privacy behaviour?*

If protection is used frequently, a number of factors might be responsible. If settings are ignored, there might be other reasons. At the end of this study, we conducted semi-structured interviews with each participant. By dissecting these discussions through PMT, we explore what drives smartwatch behaviour. This builds on our analyses in Chapters 6 and 7.

Experimental structure

Overview. The study was divided into three distinct phases: pretest, gameplay and posttest. During the 18-day pretest period, we assessed the baseline concerns and behaviour. This was undertaken for two purposes. Firstly, it captured the natural perceptions and actions of our smartwatch users. Secondly, if baseline behaviour is known, it can be compared against later actions. Through this approach, we can ascertain whether protection has changed.

After the pretest phase, all participants had a 16-day gameplay period. The sample was divided into a treatment group and a control group, seeking to limit the influence of confounding variables [344]. Group assignment was randomly determined before the study. This was undertaken in a purely-random fashion, with no stratification based on demographics. The first group received a privacy game, with its challenges dynamically customised on user behaviour. The control group were given the same app, but with the privacy customisation removed. All participants played their game over a four-day and five-day period, with a week break in the middle. This two-stage approach was selected since it can assist the refinement of mental models [257].

Finally, participants began the 18-day posttest phase. We again monitored concern and behaviour, enabling comparisons against our baseline. We also compared our groups to explore whether the privacy game might be persuasive. We completed the study with rationale interviews, capturing user experience throughout the process. An overview of this structure can be found below in Fig 8.1.

¹SQ6: *Can we mitigate the medium-term prevalence of the Privacy Paradox on smartwatches?*

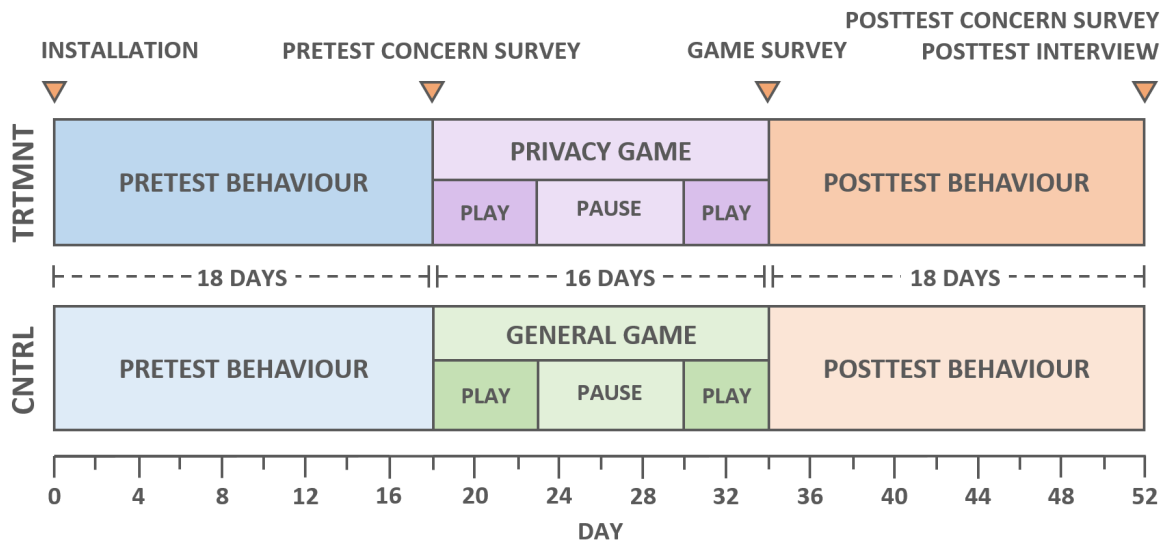


Figure 8.1: Experimental Structure

Study length. We monitored our smartwatch users over a 52-day period. Although we would have preferred a longer span, we were limited by the term lengths of our student participants. We were necessitated to recruit students, due to the conditions of our ethical approval. This is discussed in greater detail in the next subsection. We could have continued our monitoring outside of term time. However, routine actions are likely to differ during vacation periods. Therefore, to compare pretest and posttest fairly, we were constrained to an 8-week span.

Installation. The initial meetings were conducted, individually, on Day 1. 10 participants were given an Wear OS smartwatch, with the researcher assisting with installation. After configuration was finished, we installed the monitoring app on each watch. This background service logged high-level settings every five minutes. Users provided informed consent for this ethically-approved approach. Fortunately, based on interview responses, they did not appear influenced by this monitoring.

Pretest period. After the 18-day pretest phase, participants returned for another meeting. At this point, they completed a questionnaire to gauge their baseline concerns. This questionnaire is outlined later in this section. Since they had used their watches for almost three weeks, they could provide informed responses. While participants were completing the forms, we installed a game on their watches. The treatment group got the privacy game, while the control group got the generic version.

Gameplay period. Individuals were instructed to play their games three times per day for 10 minutes each time. As previously mentioned, there was a one-week pause in the centre of the period. Therefore, although the phase was 16 days, the

games were only played on nine. These games logged usage data to the smartwatch. This allowed participants' performance to be monitored.

Game survey. After this 16-day phase was completed, we distributed evaluation questionnaires. Participant responses sought to inform the refinement of future games. The forms did not concern privacy, since we wished to avoid priming the posttest period [65]. During questionnaire completion, we uninstalled the educational games. This prevented later usage, enabling a fair analysis of retention.

Posttest period. For these final 18 days, we continued to monitor privacy behaviour. At the end of this phase, we distributed privacy concern questionnaires. These were identical to the pretest forms, enabling a direct comparison. Therefore, if concerns have changed, this should be detected. This questionnaire was completed 18 days after the gameplay phase. Due to this extended period, we doubt privacy would be primed.

Interviews. Finally, to conclude the study, we conducted semi-structured interviews. These discussions explored privacy awareness (of both risk and protection), knowledge, risk perception and behavioural rationale. Through this, we analysed how well privacy knowledge was retained. We also requested justifications for changed behaviour. To avoid biasing earlier responses, these discussions were conducted after the questionnaires.

Once the interviews were complete, the smartwatches were returned to the researchers. The devices were then factory reset, preventing sensitive data from being accessed. Individuals also received their compensation: a £40 voucher and entry into a £70 draw. The study required frequent and time-consuming participation over two months. Furthermore, we sought to disguise its purpose through general advertisements and decoy questions. Finally, remuneration was allocated evenly regardless of participant responses. Therefore, we doubt this compensation influenced our findings.

Participant recruitment

Sample size. We purchased 10 Huawei Watch 2 devices for this study. We opted for this recent model since we wished our research to be future-proof. To limit extraneous variables, we monitored all our participants over the same 52-day period. Therefore, with this being single-researcher work, our sample was practically limited to 10 participants. While this is not large, comparable field studies have also had fewer than 30 individuals [22, 384]. They analysed behaviour on prevalent platforms (smartphones and social networks, respectively), whereas we explore a novel environment. Furthermore, while these studies assessed owned devices, we procured

new smartwatches. With 5-person control and treatment groups, statistical significance was largely infeasible. Therefore, we conducted a rich mixed-methods analysis, supporting quantitative findings with detailed rationale.

Sampling approach. The participants were loaned an expensive consumer device for two months. With this being university property, there was an inherent security risk. Therefore, in compliance with ethical approval requirements, we recruited from our institutions’s students. Fortunately, this demographic corresponds well with the smartwatch user population. Owners have been disproportionately young [282], with this trend appearing to continue². Smartwatch users also tend to be reasonably well-educated [112]. This should grant external validity to our student study.

Distribution. Recruitment was undertaken for four weeks between December 2017 and January 2018. For advertisement, flyers were affixed to notice boards in 25 University of Oxford colleges³. These flyers outlined the participant requirements, procedure and compensation. The study was not advertised as concerning privacy, since we did not wish to highlight our purpose. We also sent emails to mailing list curators, who could forward the messages if desired.

Screening. Individuals were invited to email the researchers if they met three conditions. Firstly, to be accountable for the smartwatch, they had to be University of Oxford students. Secondly, they should be 18 or over, so they can provide informed consent. Finally, participants were required to possess an Android smartphone. Since the smartwatch was (Android) Wear OS, this was crucial for compatibility. Most Oxford students are adults, and Android has the largest smartphone market share⁴. Therefore, we do not believe our criteria were excessively restrictive.

Informed consent. Eligible candidates were sent the participant information sheet. All participants consented to the monitoring of their smartwatch settings. We were interested in privacy metrics, such as GPS usage. However, we were concerned that knowledge of our purpose might bias behaviour [287]. Therefore, we also received consent to monitor font size, screen brightness and battery level. This disguised our privacy theme, while also ensuring high ethical standards.

Contextualised question selection

Motivation. To gauge concern, we solicited each participant’s personal opinion. As in Chapters 6 and 7, we made use of privacy scenarios. However, since we now eval-

²www.statista.com/statistics/739398/us-wearable-penetration-by-age/

³https://en.wikipedia.org/wiki/Colleges_of_the_University_of_Oxford

⁴<https://www.idc.com/promo/smartphone-market-share/os>

uated a specific device, we updated our incidents. As it was important that these questions were appropriate, we used the consistent selection criteria: *relevance*, *feasibility* and *behaviour correspondence* (Section 3.5). The *smartwatch* model (Section 3.4) was again considered when evaluating *feasibility*.

Feature selection. Before we discuss our contextualised questions, we must briefly introduce the features. This is important for the third criterion outlined above. In our prototype study (Chapter 7), we considered three smartwatch actions. These comprised: enabling a *screen lock*, *disabling GPS* and restricting *app permissions*. Their selection was justified extensively in that prior chapter. To inform our game through prototype feedback, we study the same features.

Question selection. Based on our above three criteria, we chose our concern incidents. To explore this matter in depth, we selected two questions for each feature. Through this approach, we were able to triangulate concerns to multiple issues. This resulted in three scenario pairs: *stranger access*, *location tracking* and *app access*.

Stranger access was the first pair to be considered. Concern was assessed through two incidents, both relating to this issue. In the first, we queried how users would react if their watch was accessed. Such an unauthorised invasion has clear relevance to privacy. This issue also relates to Solove’s violations [352], particularly intrusion and a breach of confidentiality. The scenario is feasible and fits within our threat model. Since smartwatches are small and valuable, they face a reasonable theft risk. Indeed, their “*size and portability makes them easy to steal*” [320]. Furthermore, they may be left unattended during sports or other activities. In these circumstances, the product could be easily accessed. Finally, there was a direct correspondence between the issue and protective features. *Screen locks* reduce the probability of unauthorised access. On smartphones, they have also been suggested as a theft deterrent [102]. Therefore, if a person is concerned about this issue, screen locks can provide protection.

In the second incident, we queried how users would react if their apps were used. This explores *stranger access* from another angle, allowing concerns to be triangulated. The issue had relevance to privacy, since strangers could impersonate the victim. It concerned the principle of appropriation, where a person’s identity is compromised [352]. This scenario was feasible, since apps are available on these smartwatches. Therefore, if an unlocked device is mislaid, the applications can be used. Finally, the risk can be reduced through *screen lock* usage. Therefore, we provided a correspondence between contextualised concern questions and features.

Location tracking was the second pair we selected. For the first incident, we asked how users would feel if their location was tracked. While our app monitored GPS

usage, we did not access the coordinates. Furthermore, participants both volunteered for the study and received compensation. Our scenario concerned company tracking, and this might place data at greater risk. The question met our three selection criteria. Firstly, it conformed to the surveillance principle [352], and is therefore relevant to privacy. Secondly, the issue is feasible, since locations are identified through GPS [180]. Smartwatches possess this feature, and it is included within our threat model. Finally, there was a correspondence between the incident and protection. If a person is concerned about monitoring, the risk can be reduced by *disabling GPS* [314].

We also queried how users would feel if their location was shared. Again, this directly related to the topic of privacy. If a private position is disseminated, it violates the disclosure and increased accessibility principles [352]. This incident is also feasible, since companies exchange information on data markets [334]. As GPS is widely used, firms have a great opportunity to collect data. Most importantly, protective actions can reduce the risk. Companies can only share the details they manage to collect. By *disabling GPS*, this quantity of data can be reduced [314].

App access comprised the final pair of questions. Firstly, we queried how users would feel if personal data was accessed. Since the scenario concerned intrusion [352], it complied with the relevance criteria. This incident is also feasible, since apps often request smartwatch data. If an app uses this information, it could be stored by its company. Finally, this scenario's risk can be reduced through feature usage. Although applications collect data, this can be limited through *app permissions*.

We then queried how users would feel if that data was shared with a third party. This incident complied with all three selection criteria. Firstly, such sharing has relevance to privacy. It relates to the 'increased availability' violation [352], as others are granted access to data. Secondly, information is frequently shared with commercial partners [334]. While details might be aggregated, they can still be placed on data exchanges [334]. Due to these facts, the incident belongs within our feasible threat model. Finally, there is a direct correspondence between risk and protective features. By restricting *app permissions*, users limit the information that can be collected. This reduces the quantity of data that could be shared. If a person is concerned about this issue, restricting permissions might be prudent.

Concern questionnaire design

Overview. Concerns were evaluated through identical questionnaires in pretest and posttest. It was important that these concerns were not inappropriately inflated. To reduce the priming risk, the form was framed as relating to general opinions.

The questionnaire could be split into three sections: *Demographics*, *Opinions* and *Concerns*. It can be found in Table 8.1, with privacy-related questions shown in bold. The queries are briefly outlined below, but discussed again in the results.

Demographics. This section appeared last, since we thought requesting data might prime privacy concerns. In previous chapters, we solicited gender, age range and highest education level. However, based on our student sample, we did not expect the latter two demographics to vary. Instead, we recorded gender, country of origin and university division. Although all participants were UK residents, we thought it interesting to explore their background. Furthermore, to encourage a varied sample, we recruited users from a range of academic disciplines. Responses did not differ by demographics, and such comparisons are omitted for brevity.

Opinions. This section sought to explore high-level viewpoints. Although four questions related to privacy, another 10 concerned other topics. This aimed to further disguise the focus of the study. Some of these queries were general in nature, enabling high-level analyses. However, six were purely decoy questions. We also included an Instructional Manipulation Check [286] to assess engagement. Respondents were tasked to provide a certain reply to indicate their attentiveness. Fortunately, all participants gave requested answers in both the pretest and posttest questionnaires. This reduces the risk that they skimmed over the form [286].

The first three queries analysed whether users enjoyed the smartwatch. Question 5 then assessed privacy awareness, checking whether users knew about data access. If consciousness increases after the game, perhaps the app has highlighted the issue. Question 8 explored our participants' understanding of privacy features. Again, through interaction with their game, we expect the treatment group to gain awareness. The surrounding queries served as decoy questions.

Finally, we tested risk perceptions, asking whether two issues were realistic. The first related to a smartwatch being lost or stolen (Q11). If an individual does not believe a threat is feasible, this might explain a lack of action. We also asked whether apps could place data at risk (Q13). Again, if practices are unknown, permissions are likely to remain open. Since our game seeks to enhance risk perception, we hope threats will become perceived.

Concerns. The contextualised questions have been introduced in the previous subsection. Again, degree of concern was solicited on a 5-point Likert Scale from 'Indifferent' to 'Very Concerned'. Based on the mean responses to each incident pair, we received three metrics. They comprised the *stranger score*, the *location score* and the *app score*. These all ranged from 1/5 (low) to 5/5 (high). In addition to

Table 8.1: Pretest/Posttest Concern Questionnaire: Privacy Questions in Bold

#	Demographics
1	With which gender do you most identify? <i>Male</i> <i>Female</i> <i>Non-Binary</i>
2	What is your country of origin?
3	In which University of Oxford division is the degree that you are studying? <i>Humanities</i> <i>MPLS</i> <i>Medical Sciences</i> <i>Social Sciences</i>
#	Opinions (Five-point Likert Scale from Strong Agreement to Strong Disagreement)
1	“I find the smartwatch useful.”
2	“I use a wide range of features on the smartwatch.”
3	“I would experience inconvenience if I didn’t use the smartwatch.”
4	“It is possible for smartwatch apps to simplify common tasks.”
5	“It is possible for smartwatch apps to access personal data.”
6	“It is possible for smartwatch apps to drain the battery.”
7	“I have a strong understanding of smartwatch notification features.”
8	“I have a strong understanding of smartwatch privacy features.”
9	It is important you remain attentive. Please mark X in the Strongly Disagree box.
10	“I have a strong understanding of smartwatch display features.”
11	“There is a realistic chance of smartwatches being lost or stolen.”
12	“If I didn’t configure my settings, my apps might drain my battery.”
13	“If I didn’t configure my settings, my apps might place my data at risk.”
14	“If I didn’t configure my settings, my apps might slow down my watch.”
#	Concerns (Five-point Likert Scale from Indifferent to Very Concerned) and Rationale
1	How would you feel if Google (the developer of Android) changed your smartwatch’s default font size?
2	How would you feel if app companies could track your precise current location?
3	Imagine a software update changed your smartwatch’s font size. How would you feel if the text was made much smaller than it was before?
4	How would you feel if app companies could read your personal data from your smartwatch?
5	Imagine your smartwatch was lost or stolen. How would you feel if a random stranger could read your data?
6	How would you feel if Google (the developer of Android) changed your smartwatch’s default alarm volume?
7	Imagine your smartwatch was lost or stolen. How would you feel if a random stranger could use your apps as you?
8	Imagine a software update changed your smartwatch’s alarm volume. How would you feel if the alarm volume was set much quieter than it was before?
9	How would you feel if app companies could share your personal data with other companies?
10	How would you feel if Google (the developer of Android) changed your smartwatch’s default screen brightness?
11	How would you feel if app companies could share your precise movements with other companies?
12	Imagine a software update changed your smartwatch’s screen brightness. How would you feel if the brightness was set much lower than it was before?

the discrete responses, participants provided a textual justification for each answer. These remarks provided a deeper insight into smartwatch concerns.

Since we wished to disguise the topic of privacy, we again included decoy queries. To achieve this, we created another three pairs of contextualised questions. These were syntactically-similar and related to general smartwatch usage. The pairs concerned: *font changes*, *volume changes* and *brightness changes*. However, if questions bore no relationship to the game, they might be identified as decoys. Therefore, the control-group challenges were given a correspondence with these scenarios. For example, one incident considered the screen brightness being reduced too low. To match this, the (control-group) game asked users to adjust their brightness. Participants also believed these details were logged, as it was approved at the start of the study. These approaches should help disguise the study's purpose.

Evaluation questionnaire design

Summary. Both groups played a game during the gameplay period. To inform future apps, we solicited feedback through a brief questionnaire. This form also allowed us to evaluate the success of our implementations. If they are considered usable and enjoyable, it suggests they have been well-designed.

The final form was split into two sections: *Evaluations* and *Opinions*. Our queries can be found in Table 8.2. For brevity, they are briefly introduced and then described in the gameplay results (Section 8.5). The former section sought to explore whether games were considered enjoyable, usable and educational. It also evaluated game difficulty, allowing the apps to be compared. The latter section solicited which elements were most and least appreciated. This aimed to inform the development of future games. It also assessed retention and response veracity. The questionnaire did not concern the topic of privacy. Since it immediately preceded the posttest phase, we did not wish to prime the concept.

Posttest interview design

Purpose. While we assessed questionnaires and logs, rationale is challenging to extract from these sources. Therefore, we concluded the study by conducting semi-structured interviews. We selected this format as it supports both exploration and the comparison of responses [285]. Since these discussions took place on the final day (Day 52), participants should be able to contribute informed opinions.

Table 8.2: Game Evaluation Questionnaire

#	Evaluations (Five-point Likert Scale from Strongly Agree to Strongly Disagree)
1	“I found the smartwatch game to be enjoyable.”
2	“I found the smartwatch game to be usable.”
3	“I found the smartwatch game to be educational.”
4	“The game taught me things that I did not previously know about the smartwatch.”
5	“I found the challenges in the smartwatch game to be easy.”
#	Opinions (Qualitative responses)
1	Approximately, what was the highest score that you received?
2	What did you like most about the smartwatch game? Why?
3	What did you like least about the smartwatch game? Why?
4	What about the game would you like to see improved? Why?
5	What different colours were the houses in the smartwatch game?

Procedure. The 40-minute interviews were conducted in a university meeting room. The discussions were one-on-one, reducing the influence from other participants. To support a rich qualitative analysis, the session was audio recorded. Participants gave their explicit consent and were able to suspend the recording at any time. They were also free to retract existing comments, although this never occurred.

Structure. The interview had four sections: *General*, *Awareness/knowledge*, *Protection Motivation Theory* and the *Privacy Paradox*. The questions can be found below in Table 8.3. We outline our queries, though they are described in Section 8.7.

General. This section sought to capture high-level perceptions of the study. The questions also avoided the topic of privacy, since we did not wish to bias later responses. We first solicited each user’s experience in wearing the smartwatch. We then asked them why they participated, assessing their motivation. If they were driven by curiosity rather than compensation, we might have greater external validity. Bravely, we continued by asking whether the background app affected their actions. If not, this may grant credibility to our findings. Finally, we explored retention by soliciting whether anything had been learned. We expect our treatment group to have gained knowledge of protection.

Awareness/knowledge. In this section, we sought to explore the influence of the games. While individual differences will exist, we expected the treatment group to have gained greater knowledge. Questions 6/7 concerned PMT, but were asked at this stage to reduce order effects. To evaluate awareness, we asked users how their data might be accessed by others. Since awareness might not be sufficient for behaviour change [38], we continued by exploring knowledge. Participants were asked what they

Table 8.3: Posttest Interview Questions

#	General
1	What was your experience in wearing the smartwatch?
2	Why did you choose to participate in the study?
3	Do you feel the background StudyService app affected your behaviour? Why?
4	Would you purchase your own smartwatch? Why?
5	Do you feel you learned anything new as the study progressed? If so, what?
#	Awareness/knowledge
6	How likely do you believe the chance of companies accessing your watch's data? Why?
7	How likely do you believe the chance of someone's smartwatch being lost or stolen? Why?
8	How privacy-conscious do you generally consider yourself to be? Why?
9	How do you think your smartwatch's data could be accessed by companies or other people?
10	Imagine your smartwatch settings were changed back to their defaults. If you wanted to, what could you do to protect your smartwatch's data? Why?
11	Imagine your smartwatch settings were changed back to their defaults. If you wanted to prevent apps from tracking your location, what could you do? Why?
12	Imagine your smartwatch settings were changed back to their defaults. If you wanted to stop apps from reading your personal data, what could you do? Why?
13	Imagine your smartwatch settings were changed back to their defaults. If you wanted to limit watch access in case of loss or theft, what could you do? Why?
14	Could you please show me, and explain aloud, how to disable GPS on your smartwatch?
15	Could you please show me, and explain aloud, how to change the permissions for a smartwatch app?
16	Could you please show me, and explain aloud, how to enable a screen lock on your smartwatch?
#	Protection Motivation Theory
17	On a scale from 1 (low) to 10 (high), how serious do you feel the action of your smartwatch data being accessed by a company is? Why?
18	How effective do you think smartwatch settings can be in protecting your device's data? Why?
19	How able do you feel you are to protect your smartwatch's data? Why?
20	Do you feel you receive benefits from using data-accessing apps? If so, what?
21	How much effort do you feel it is to protect your smartwatch's data? Why?
#	Privacy Paradox
22	We have discussed the use of tools which protect your smartwatch's privacy. Can you think of any techniques or circumstances that would lead you to use these tools more often?
23	Most of us claim to be concerned about our privacy. However, most of us also fail to fully protect ourselves. This contrast is known as the Privacy Paradox. Why do you think this situation might occur?
24	You have indicated that you are concerned about your smartwatch's data being accessed. However, on occasions, you didn't use settings to protect that data. Why do you feel this was the case?
25	Is there anything else you would like to add that we haven't discussed?

would do to protect their data. They were told to respond as if their settings were at their defaults. This avoided empty responses when the devices were already guarded.

To further assess their understanding, we probed their responses to defined issues. To support a correspondence with concern, users were asked how they would limit: *stranger access*, *location tracking* and *app access*. If the treatment group identified the privacy features, their game might have taught protection. As the control group played the generic version, we expected their baseline knowledge to be low. However, even if a person knows of a feature, they might be unable to use it. Therefore, we explored genuine knowledge by requesting a demonstration. Participants were tasked to: enable a screen lock; disable their GPS; and change their permissions. They used their smartwatch and explained aloud for the audio recorder. If treatment participants can master these settings, their game might be educational.

Protection Motivation Theory. This section dissected behavioural rationale, seeking to understand user decisions. Establishing consistency with prior chapters, we based questions on Protection Motivation Theory (PMT) [328]. We examined each factor through its own question. For vulnerability, we asked users how likely it was that data was accessible. If individuals do not perceive a risk, this might explain a lack of protection. To evaluate severity, we requested a rating of the seriousness of data access. We cared little about the quantitative figure; it simply served to trigger discussion. This approach was preferred to mentioning ‘severity’, which might have encouraged negative responses. If a violation is deemed to be damaging, individuals should be more likely to protect themselves.

We then explored response efficacy by asking how effective settings were considered to be. We also assessed self-efficacy, by querying each participant’s degree of confidence. If individuals doubt protection or their usage of it, they might neglect privacy features. For rewards, we asked users whether they benefited from data-accessing apps. When apps provide useful functionality, information might be disclosed. Finally, to assess response costs, we asked how much effort it is to protect data. If settings are opaque or complex, users may be less likely to configure them. Through exploring the impact of these factors, we can understand why protection is (not) used.

We chose not to analyse these components merely through response proportions. We sought to develop ‘participant profiles’ for each individual. These would comprise of single-paragraph descriptions of a person’s rationale. Since their accuracy is important, the content will be verified through respondent validation. Based on these rich profiles, we can assess how participants make privacy decisions.

Privacy Paradox. The final section directly explores our thesis topic. It seeks to understand why the Paradox might exist in a smartwatch environment. We begin by asking which circumstances would encourage greater protection. However, since responses concerned environmental changes (e.g., new job, moving countries), they were not applicable to our techniques. We continued by requesting a justification for the Paradox. This takes a similar approach to our Chapter 6 interviews, but focuses on smartwatches. If several individuals reference a certain factor, it might be influential in wearable environments. Now that the topic was introduced, we suggested to participants that their actions did not always align with concerns. To avoid antagonism and defensive responses, this issue was explained to be common. Through their own justifications of behaviour, we extract insights into user rationale. This enabled a direct exploration of the Paradox in a smartwatch environment.

8.3 Smartwatch Game Design

We now continue by discussing the design of our two smartwatch games. Since our privacy version sought to encourage protection, most attention will be given to this app. However, aside from the privacy customisation, the games were identical. Screenshots from the privacy version can be found below in Figure 8.2. The leftmost image displays the game map, while the rightmost presents a privacy challenge. For brevity, additional images are located in Appendix B.2.

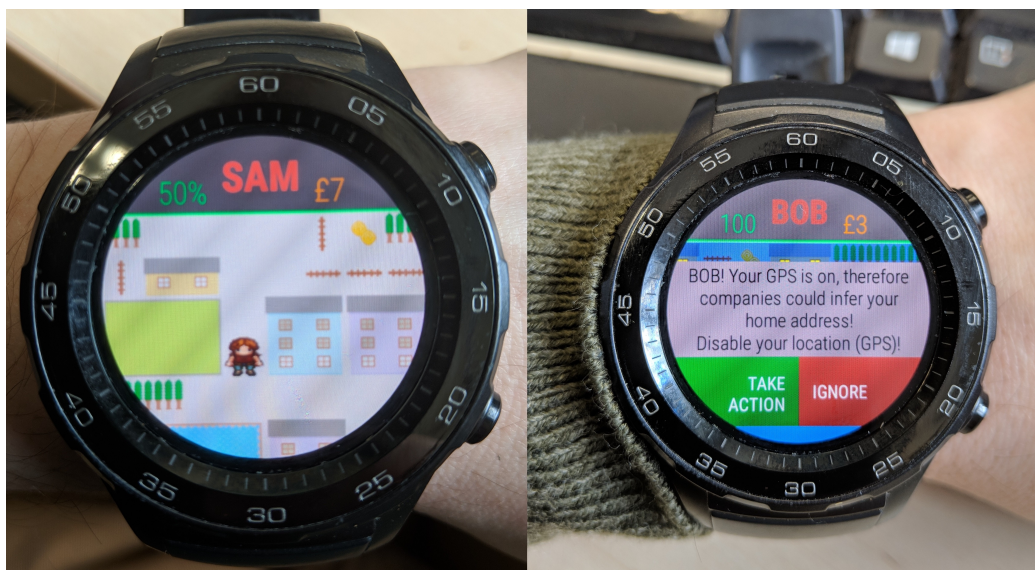


Figure 8.2: Smartwatch Privacy Game: Gameplay (left), GPS Challenge (right)

A video of the game can be found at www.youtube.com/watch?v=1DZ29t0_z0Q. In addition to these two programs, we also developed the background monitoring app. Since behavioural analysis is central to our study, we begin by describing this app.

Background monitoring app

Monitored data. As previously mentioned, settings were logged once every five minutes. These details, first introduced in the Research Questions subsection, correspond with the privacy-protective behaviour. Accordingly, we logged whether a screen lock (pattern, PIN or password) was on. We also recorded whether GPS access was currently enabled. Finally, we collected the list of apps and their permissions. Through this data, we monitored how behaviour varied over 52 days. Although we were approved to collect additional non-privacy details, this was deemed unnecessary.

Frequency. The background application was required to function consistently over a two-month period. It made use of the Android AlarmService, allowing logs to be produced even after the app was garbage-collected. Wear OS has several battery-saving procedures, one of which is grouping alarms. Therefore, the service was not guaranteed to run at exact five-minute intervals. However, the interval never exceeded 10 minutes and was usually less than five. Even if the gap was 10 minutes, the 144 records per day would provide sufficient monitoring.

Interpolation. Other constraints were created by Wear OS battery procedures. Services do not trigger when the watch is idle for long periods. However, the act of adjusting settings (or merely touching the screen) would wake the interface. Even updates or USB access would cause this awakening. Therefore, while the watch was in an idle state, settings were identical to the previous log.

Android updates. Halfway through the study, Google updated the Wear OS operating system. The new version introduced another battery-saving procedure. If the watch was idle for long periods, GPS was disabled to conserve power. This was not behaviour undertaken by the participant, but an unprompted background alteration. Our research sought to explore privacy-protective actions, rather than involuntary changes. After the study was completed, we consulted with our participants. All individuals confirmed that they did not disable GPS during these periods. Indeed, through interview demonstrations, we saw that many were unsure where settings were placed. Since we study behaviour, we disregarded incidents of involuntary disabling.

Game narrative and mechanics

Summary. We now introduce our two educational games. To minimise confounding variables, both were entitled ‘Shop Dash’ and concerned the same theme. The narrative is similar to that of the Chapter 7 app, since it was informed by that application. Furthermore, the concepts of questions and challenges are also retained. The same Learning Science principles were adopted for both games. To reduce repetition, this chapter only outlines the extensions from the online prototype. For example, the final app includes an introductory slideshow when it is first opened. This highlights privacy risk and demonstrates how to use protective settings.

Questions. As before, users navigate their character from its home to the local shops. Unlike the previous game, they must now traverse four levels to reach this point. During their journey, users encounter two varieties of Non-Player Character (NPC). The first, ‘villagers’, ask functionality questions and reward correct answers. Since queries related to the application theme, privacy questions concerned the protective features. In contrast, the control group are asked about general settings. Through refining knowledge of the device, we hope to provide education.

Challenges. The second NPC, ‘thieves’, differ in that they trigger functionality challenges. As in the prototype, challenges require users to adjust a watch’s settings. In the privacy game, they concerned the three protective features. Therefore, participants gained direct experience in protecting their watch. Through learning how to guard data, concerns might also be moderated. The control-group game assigned general tasks, such as adjusting screen brightness. This provided a correspondence with the questionnaire’s decoy queries.

Progression. When the player reaches the end of each map, they move to the next level. At this point, their challenge and question success is displayed. Reflection often increases retention [120], particularly when users consider their own progress. After traversing all four levels and reaching the shops, they complete the game. Additional difficulties are then unlocked, with tougher challenges/questions but greater points. To climb the leaderboard, users are incentivised to refine their skills.

Game challenges

Interface. In the privacy game, the challenges aligned with the protective features. Rather than adjusting the real settings menu, we implemented a simulated interface. The interface was identical to the real dialogues, but provided a number of advantages. It allowed navigation to be monitored and helpful hints to be displayed. It also meant

that individuals could experiment in a safe environment. If they were tasked to change real settings, it would have impacted our behavioural monitoring.

Distribution. In both games, participants faced a total of 14 challenges. Their difficulty increased with the levels, with quicker settings allocated in the earlier stages. When additional difficulty modes were unlocked, the challenges remained the same. However, health decreased quicker, encouraging memorisation of menus.

Privacy challenges concerned the three protective features. By using multiple tasks per feature, we sought to refine knowledge. These tasks can be found below in Table 8.4. The set challenges are in italics, while the others were randomly allocated. This randomisation sought to keep the game unpredictable and fun.

Table 8.4: Privacy Game Challenges

Level One (2 Challenges: 1 Set, 1/3 Random)	
<i>Disable GPS</i>	Enable a screen lock pattern
Check app permissions	Revoke contacts permissions
Level Two (3 Challenges: 1 Set, 2/3 Remaining from Before)	
Enable a screen lock pattern	Check app permissions
Revoke contacts permissions	<i>Revoke audio permissions</i>
Level Three (4 Challenges: 1 Set, 3/3 Random Order)	
Enable a screen lock PIN	Revoke location permissions
Check system app permissions	<i>Disable GPS and location permissions</i>
Level Four (5 Challenges: 5/5 Set Order)	
<i>Revoke SMS permissions</i>	<i>Enable a screen lock password</i>
<i>Revoke sensors permissions</i>	<i>Enable pattern, disable GPS and location perms</i>
	<i>Uninstall application</i>

Dynamic risk exposure. In Chapter 7, participants suggested they would be influenced by risk exposure. If true, this implies that protection might be encouraged by highlighting threats. To communicate this risk, privacy challenges were customised around participant behaviour. Rather than a task presenting generic text, the issues were personalised to each individual.

This was achieved through reading the recent log files of the monitoring app. For example, a participant might have granted ACCESS_FINE_LOCATION permissions to their Uber application. If they subsequently faced a location challenge, it would be contextualised with this information. This followed the influential approach of Harbach et al. [165], who highlighted risk exposure on smartphones. To encourage protection, we sought to establish a clear correspondence between behaviour and consequence. This contextualisation was dynamic, updating based on recent recordings.

Therefore, the challenges adjusted in response to changes in behaviour. This dynamic feedback loop aimed to persuade our participants.

Contextualisation. Understanding is enhanced when physical and virtual risks are aligned [150]. Therefore, we matched challenges to real-world scenarios. For example, characters face a challenge when they are near their (gameplay) house. This task requires GPS disabling, since their ‘home’ location is being tracked. On another occasion, a thief and villager are adjacent to each other. Since questions could be overheard, the challenge concerns microphone eavesdropping. By relating risks to real situations, participants might consider threats in the future.

If challenges were failed, potential consequences were highlighted. For example, a participant might fail a GPS task while Uber is installed. In this case, we explained the possibility of location tracking. Therefore, we sought to further demonstrate a behaviour-consequence connection.

Performance monitoring. Participants were required to play the games three times per day for 10 minutes each time. To both monitor compliance and analyse performance, we monitored gameplay activity. The app logged name, avatar and difficulty when a game was started. When the games finished, we recorded level, total score and challenge results. By comparing performance across participants, we could identify which topics posed the most difficulty.

Protection Motivation Theory

Purpose. In this thesis, privacy is deconstructed through PMT [328]. By considering perceived risks and response costs, individuals decide whether to guard their data. Since we wish to encourage protection, we implemented PMT factors in our game.

Approach. If a person does not consider themselves vulnerable, they are unlikely to act. Therefore, we highlighted participants’ risk exposure through the personalised challenges. Similarly, they might reject protection if invasion is deemed not severe. Through the introductory slideshow, we explained the potential consequences.

Self-efficacy might be lacking when participants begin the study. However, by granting opportunities to master the features, confidence should increase. Response efficacy should also be enhanced, since the slideshow highlights the advantages of settings. Therefore, by presenting the threat and supporting the response, protective behaviour should be encouraged.

Informed refinements

Through our large-scale prototype (Chapter 7), we received an abundance of feedback. This data informed refinements to our smartwatch implementation.

Appearance. The general theme of the game was appreciated. It was simple to understand and the narrative was accessible. Therefore, the shopping topic was retained. The game graphics were more contentious, dividing the opinion of online participants. Some named the design as their favourite element (5.6%). However, others criticised the appearance (11.9%). Based on this feedback, we retained the graphics while refining their design.

Game length. Some respondents believed the prototype was too short. To both add variety and increase content, we extended the games. Firstly, we designed and implemented another three maps. Secondly, we increased the number of questions from three to eight. Questions differed across the three difficulties, contributing to a total of 24 queries. Thirdly, we increased the challenge count from three to 14. This was achieved by combining and diversifying tasks. Therefore, even if the game was completed only once, users would have experience in configuring settings.

Randomisation. Since the prototype was played briefly by each participant, there was no randomness. The thieves were placed in constant locations and triggered the same challenges. This sought to establish a consistent experience for the respondents. As individuals now played the game frequently, we introduced greater variety. Thieves appeared in random locations and triggered random tasks. Question order was also randomly shuffled, further adding variety. This should reduce the repetitiveness of the game, hence improving the user experience.

Feedback. Finally, we wished to give participants greater support in completing challenges. If a person is unable to locate their settings, they might lose motivation. Therefore, we used several techniques to aid their education. For example, when health decreased past 50%, helpful hints were provided. This assisted participants in navigating to the feature. If the challenge was failed, the correct route was then described. Through this, the individual should have greater success on their next attempt. By supporting participants, protective approaches should be memorable.

8.4 Baseline Findings

Participants

We recruited 10 participants, each of which used a smartwatch for two months. Eight of these individuals were male, while the other two were female. Since the University of Oxford is highly multinational, countries of origin were diverse. While four came from the UK, we had single participants from Ireland, Italy, Russia, Mexico, Singapore and the US. However, all were resident to the UK and completed the study in this country. This was important for our exploration of the Paradox. Academic disciplines were also varied. Five were from Mathematical Sciences, two did Humanities, two studied Social Sciences and one did a Medical Science. Rather than evaluating technology experts, this provided a more-representative sample. As previously mentioned, the smartwatch user population has been disproportionately male and young [282]. Therefore, our sample should possess decent external validity.

Qualitative techniques

Coding frame creation. Through our three questionnaires and our interviews, we collected a large quantity of qualitative data. To support a robust evaluation, we again undertook inductive analysis [323]. This approach is described in detail in Section 3.3. We used this process for all our qualitative data, enhancing the validity of our findings. As a result, we received 40 structured indices. 24 of these came from the interviews, 12 from the concern questionnaires and 4 from the evaluation forms. They were each refined into coding frames, supporting our rich analysis. Those frames relevant to our textual discussion can be found in Appendix A.4.

Deviant cases. When coding responses, themes can differ greatly in proportions. Since many remarks might be mentioned once, it can be tempting to subsume these in other categories. However, we respected the deviant cases [23]. Our concern questionnaires had several such incidents. In Question 4, participants reported their response to app data collection. Only one person was concerned about ‘Identification’, but this was not subsumed into the ‘Personal data’ theme. Interview responses were also highly varied. In Question 1, individuals were asked for their general experience. One person found the watch large on their wrist, but this was not subsumed within usability issues. By respecting such cases, we sought to remain faithful to our participants’ responses.

Coding approach. Distinct frames were constructed for each of our qualitative questions. This allowed us to analyse participants from a range of angles. When evaluating concerns, individuals completed both the pretest and posttest questionnaire. To compare responses across the study, their frames were combined. Therefore, one table was used for Q1 responses, whether analysing pretest or posttest. By exploring how proportions vary, we could assess whether concerns had changed.

As in our previous studies, a person's response could mention a number of topics. Because of this, we might receive more codings than the number of participants. Therefore, when calculating proportions, we use the 'comment' format (Section 3.3).

Qualitative validation

To add further rigour to our process, we then performed qualitative validation. We made use of all four techniques outlined in Section 3.3. These were: *triangulation* [139], *repeated coding* [245], *multiple coding* [293] and *respondent validation* [56]. Although the approaches are previously described, we briefly reiterate the process. We also outline the results in each case.

Triangulation. In this study, we receive results through multiple sources. By comparing these results, we can corroborate our findings [139]. We triangulated through using both questionnaires and interviews. In the concern surveys, we explored participants' reactions to privacy incidents. This is undertaken in both pretest and posttest. Then, through the interviews, we also discuss the potential risks. Since rationale aligned in both cases, we have confidence in these findings.

Behaviour was also explored both empirically and through discussion. While the logs indicated how participants acted, it did not provide their rationale. However, behaviour was also examined through the interviews. Knowledge could be assessed based on the answers to privacy questions. Those lacking knowledge also tended to neglect their protection. By using these two instruments, behaviour was corroborated.

Repeated coding. We sought for our coding process to be as rigorous as possible. This was undertaken by having refined frames and strong definitions. However, when data is coded once, subjectivity might have some influence [346]. Therefore, through repeated coding, we analysed our *intra-rater* reliability [245]. Again, the exact details are outlined in Section 3.3. This technique was used for all our qualitative data, whether from concern questionnaires, evaluation forms or final interviews. By comparing our coding results, we could identify areas of ambiguity. We could then reconcile our frames and strengthen our definitions.

Therefore, the doctoral researcher coded all the data twice. Once the first parse was complete, we populated a framework matrix. As described in Section 3.3, columns denoted participants and rows represented questions. Matrices were created for the concern questionnaires (pretest and posttest), the evaluation forms and the interviews. Three weeks later, we performed the second coding. At this point, another matrix was populated with the updated results. By studying the ‘proportion agreement’ metric [271], we analysed the consistency of our process. As previously highlighted, this metric was favoured over Cohen’s kappa [93] for two reasons. Firstly, there were a large number of themes, reducing the risk that matching is due to chance. Secondly, since responses often mentioned multiple themes, kappa was not appropriate [93].

Repeated coding: Results. In the pretest and posttest questionnaires, participants answered twelve *concern* questions. Six of these were decoy and therefore not considered. Across 10 participants in 6 pretest and posttest queries, we received 120 responses. Since users expressed broad rationale, responses often related to multiple codes. Across the two rounds, codes were distinctly assigned on 304 occasions. Of these, 9 were unique to the first round, while 2 were unique to the second. This contributed to an agreement rate of 96.4%. As the vast majority of codings were consistent, we believe our frames are robust. After reflecting on the results, we then selected the most appropriate themes for the 11 mismatches. Finally, the definitions were updated to mitigate future ambiguity.

We also undertook repeated coding of the evaluation questionnaires. These forms requested three pieces of qualitative information. Across our 10 participants, this resulted in 30 responses. 54 themes were distinctly assigned across both rounds, with only one assignment unique to the first. Therefore, we received an agreement rate of 98.1%. This high consistency gives us confidence in our evaluation results. After reflection, the mismatch was categorised and our definitions were updated.

Finally, we conducted a second parse of the interview transcripts. Each participant answered 25 questions, resulting in a 250-response dataset. As the replies were conversational, multiple themes were mentioned in most responses. This led to 687 distinct assignments across both rounds. 51 were unique to the first round, while only 28 were unique to the second. 608 were identical, producing an agreement rate of 88.5%. Although this is lower than the other percentages, this was expected. Due to the conversational nature of interviews, often six codes applied to a single response. Therefore, it could be difficult to identify all these themes consistently. Considering this challenges, the matching accuracy was strong. Finally, we reflected on the mis-

matches and updated the definitions. Through this repeated coding, we believe we strengthened our frames.

Multiple coding. Researchers can often judge a matter in different ways. Provided with the same data and coding frames, individuals might produce a variety of distributions. But if the frames are highly precise, the results can be more consistent. To introduce further rigour into our analysis, we undertook multiple coding of our posttest interviews [293]. The approach is outlined in Section 3.3. By comparing another researcher's results with our own, we analysed the robustness of our frames. While this researcher was a member of our team, they had recently joined the university. They were unfamiliar with the student's work and did not have a Computer Science background. We thought this inexperience would test our frames to the greatest extent.

Multiple coding: Results. As mentioned in Section 3.3, we populated two framework matrices. The first encapsulated the results of our repeated coding. The second was based on the coding of the independent researcher. Since 10 participants answered 25 questions, we received 250 responses. However, as interviews are conversational, 672 distinct assignments were made across both parses. Of these, 162 were unique to the first parse, while 55 were unique to the second. The other 455 codings matched exactly. This contributed to a proportion agreement rate of 67.7%. As shown, most of the unique assignments were made by the doctoral student. Indeed, this provided 74.7% of the mismatches. On reviewing the process, the second coder appeared to miss many of the themes. This is unsurprising, since some responses related to six topics. On reflection, a security academic might have been preferable to a new researcher. As a result, our repeated coding provided an analysis with greater richness. Indeed, when reconciling the matrices, the first parse was deemed richer on most occasions. Where disagreement was common, we further strengthened our definitions. Since our themes were refined through repeated coding and multiple coding, we believe the frames are robust. As demonstrated below, participants were in agreement with all the theme assignments.

Respondent validation. To verify that we encapsulated our participants' views, we involved them in the final process. Therefore, we undertook a process of respondent validation [56]. The exact approach is outlined in Section 3.3. Through this, our participants were assigned three tasks. Firstly, they verified that the transcripts were accurate. Secondly, they judged whether their assigned codes were appropriate for their responses. Finally, they verified their brief participant profiles. These profiles were one-paragraph characterisations, developed based on their responses to PMT

questions. The descriptions sought to encapsulate the behavioural rationale of each user. By checking this content, individuals verified that the profiles were fitting.

Respondent validation: Results. Transcripts were sent to all 10 of our participants. Fortunately, all 10 replied and took part in the validation process. Across the entire sample, not a single amendment was requested. All respondents verified that the transcript was accurate and the codes were appropriate. They also believed that their profile encapsulated their rationale. This 100% accuracy rate suggests that we captured our participants' views. It also gives us confidence in our final coding frames. Based on this qualitative data, we explore smartwatch behavioural rationale.

Participant opinions

Introduction. Our pretest concern questionnaires were completed 18 days into the study. Participants should then be familiar with their devices, but not in possession of training. Although we hesitate from generalisation, this might be a similar state to many smartwatch users. Since neither group has played a game at this stage, we discuss findings in terms of the whole sample.

While our questionnaires gauged concern, they first explored general opinions. While privacy elements were included, most of the questions were decoy. Therefore, we believe our topic of study was adequately disguised. The Opinions section could be subdivided into five themes: *attention*, *satisfaction*, *awareness*, *understanding* and *risk perception*. The pretest views are outlined below.

Attention. We wished to verify that our participants were attentive. To achieve this, we included an Instructional Manipulation Check [286] halfway through the questions. All respondents gave the correct reply, suggesting they invested sufficient effort. Therefore, we have greater confidence in the questionnaire's results [286].

Satisfaction. The devices were considered to be useful, with 90% agreeing with this statement. However, since none were in strong agreement, the benefits might not have been extensive. We then assessed whether participants used a range of functionality. 70% gave a neutral response, implying that the products had been moderately investigated. Finally, we asked whether the devices' removal would cause inconvenience. Since 70% disagreed with this statement, the watches appear to be far from a necessity.

Awareness. To assess this topic, we asked users whether personal data could be accessed by apps. An informed individual should know this is possible, and fortunately all participants agreed. Since 60% were in strong agreement, our sample has

decent baseline awareness. If users are aware of the risk, they should be more likely to use protection.

Understanding. Through asking individuals to rate their understanding, we also assess their self-efficacy. At this pretest stage, only 40% claimed to possess knowledge. This might discourage participants from using protective features.

Risk perception. If a person does not perceive threats, they might lack the motivation to act. To explore this matter, we first asked whether smartwatches may be lost or stolen. Since these devices are small and valuable, they should face some risk [40]. We received a cautious response, with 30% in agreement and 40% in disagreement. We then asked whether an app might threaten smartwatch data. In this case, only 40% believed the threat was possible. Since our privacy game seeks to highlight risk, these percentages should increase.

Concerns and behaviour: Lock screens

Introduction. As previously mentioned, concerns were gauged through reactions to six contextualised questions. Responses were given on a five-point Likert Scale from ‘Indifferent’ (1/5) to ‘Very Concerned’ (5/5). By taking a mean of each pair, we received the *stranger score*, the *location score* and the *app score*. A graphic illustrating the reactions can be found in Figure 8.3. To investigate the root of these concerns, we also explore the qualitative justifications for each answer. These opinions can then be compared against the records of pretest behaviour.

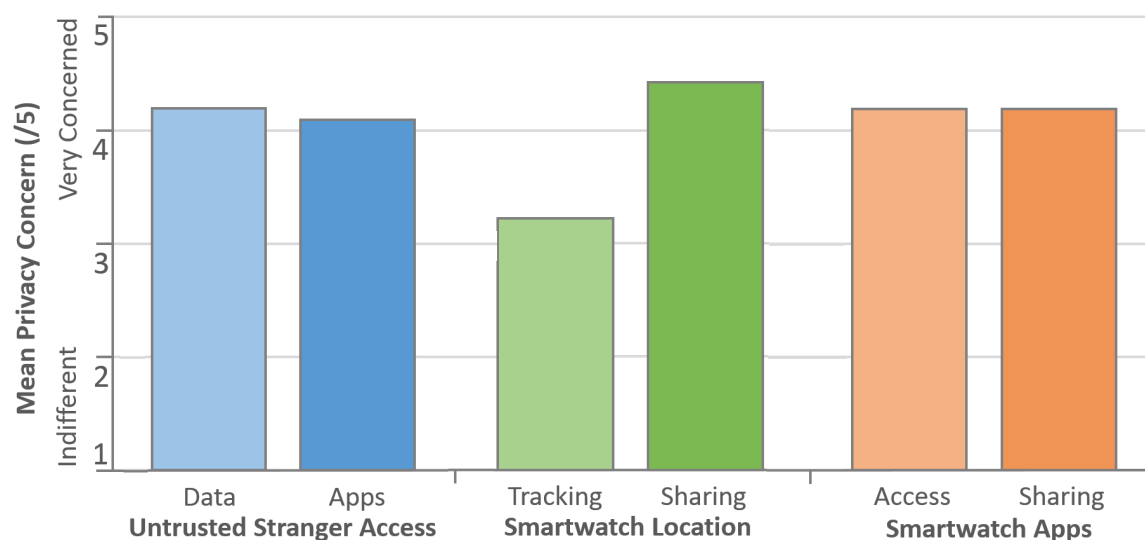


Figure 8.3: Pretest Smartwatch Privacy Concerns

Concerns. *Stranger access* was the first incident pair. When considering data reading, participants appeared highly concerned. 50% indicated their strong opposition, while no respondents were indifferent. This contributed to a mean score of 4.2/5, suggesting users reject this intrusion. When assessing these scenarios, we also considered the participants' rationale. 15 justifications, were given, with 60% of these fearing great damage. Concerns were primarily driven by the security risk (20%) and the importance of personal data (20%). This suggests that our users place some value on their smartwatches. Two individuals were less concerned about this potential issue. One believed that that their screen lock would provide adequate protection. If protective features are enabled, the risk might indeed be mitigated. Representative quotes, along with the participant ID, are shown below:

“I don't want a stranger know my usual whereabouts”, (#G).

Participants expressed similar concerns over their apps being used ($\bar{x} = 4.1/5$). 70% of users were in opposition, with only one person expressing indifference (10%). This proportion suggests that application access is strongly rejected. On this occasion, 14 comments were provided for justification. 64.3% of these feared damage, with data access being the most common concern (28.6%). These individuals were worried since apps can contain sensitive details. This might lead to financial crimes or personal danger. Only one participant was unconcerned, and they believed they possessed no 'sensitive apps' (7.1%). When checking the behaviour logs, they indeed had no applications accessing private details. However, when app usage is feared, a screen lock might be appropriate.

“They could cause issues through contact and they could gain my details”, (#E).

Behaviour. Based on the two concern incidents, we calculated a *stranger score* of 4.05/5. It appears our participants are concerned about unauthorised interactions. To explore empirical actions, we analysed our smartwatch logs. Through assessing baseline behaviour, we can address this chapter's first research question (RQ1). To reiterate, this explored whether smartwatch users take action to protect their data.

Fortunately, four participants had enabled a screen lock. Since they did not use it initially, this contributed to a behaviour *lock score* of 38.6%. Games had not been played at this point, which suggests that the feature is well-known. This might be due to the prevalence of smartphone PINs and patterns. The other six participants never used this lock. As in our Chapter 6 and 7 findings, it appears password usage is

far from constant. These inactive users still claimed concern over the scenarios ($\bar{x} = 4.2/5$). In these situations, it is possible that locks were never noticed. Considering RQ1, this implies that protective action is far from constant. RQ2 explores the baseline Paradox, and behaviour appears often misaligned with concern.

Concerns and behaviour: GPS usage

Concerns. The second incident pair was that of *location tracking*. To assess the matter, we first solicited reactions to location monitoring. Our participants appeared moderately concerned, with 50% giving negative responses. However, opinions varied, since another 30% seemed reasonably indifferent. This contributed to a mean of 3.2/5, suggesting that this tracking is not highly feared. Concerns might be mitigated by the advantages that GPS can offer. When considering the justifications, 16 comments were made. While 37.5% expressed concerns, 37.5% were contingent on the particular situation. For example, 18.8% claimed their reaction depended on whether the tracking was optional. If it could be disabled, as GPS can be, they would be less worried. However, in order to protect data, intentions must turn into actions.

“I want to be able to decide when I can be tracked or not”, (#C).

While tracking triggered the least concern, location sharing faced strong opposition. For this incident, 80% were concerned and nobody expressed indifference. This resulted in a high mean of 4.4/5. Although participants accepted company tracking, they wanted the data protected from others. We then considered justifications, with 21 comments provided. The vast majority expressed concern (85.7%), with most individuals objecting to the principle (23.8%). Participants also thought this was illegal (14.3%), although it might be consented through privacy policies. If individuals are opposed to this sharing, access can be limited by disabling GPS [314].

“I would very much feel as though my privacy was invaded”, (#A).

Behaviour. When calculating the concern *location score*, participants received 3.8/5. This balanced metric is based on tracking ambivalence and sharing concern. Individuals did not reject monitoring, provided it was optional and conspicuous. This might be due to the normalisation of digital surveillance [353]. However, the tracking capabilities of GPS might not be recognised. We continued by assessing our participants’ empirical behaviour. At this pretest stage, all 10 had the service enabled. Indeed, not a single person had adjusted this setting. Because of this, their location

was potentially accessible for the first 18 days. This contributed to a *GPS score* of 100%. This is comparable with the rare disabling found in our prototype study. When considering RQ1, it appears that protective actions are infrequent. This presents a worrying baseline for untrained users. Although concerns were not extreme, respondents did express some opposition. Little was done to mitigate the scenario risks. Menfors and Fernstedt [263] found similar results in their Paradox study on geotagging. To address this issue, our privacy game seeks to encourage protection.

Concerns and behaviour: App permissions

Concerns. Finally, our participants considered the *app access* issues. The first incident concerned an app accessing personal data. Respondents were strongly in opposition, with 90% disliking this situation. One user appeared less concerned, and their rationale will be discussed shortly. These responses contributed to a strong mean of 4.2/5. When considering the justifications, 16 different comments were given. While 25% were contingent on particular details, 62.5% expressed strong concern. Many objected to the principle of access (25%), whether or not it placed them at risk. This sentiment was also frequently expressed during the interviews in Chapter 6. This might imply that unease does not require a severe violation [303]. The indifferent participant believed that data would be accessed in other ways. If individuals are fatalistic, they might regard protection as pointless [357].

“They already can with phone, or at least I assume so”, (#A).

For the final contextualised question, we solicited reactions to app data sharing. 80% opposed this incident, while no individuals were indifferent. This resulted in a mean of 4.2/5, same as in the access scenario. Our participants seem to fear these app practices, despite them being commonly found [334]. This might imply that users have little awareness of the procedures [99]. We then considered the qualitative justifications, with 17 comments provided. 58.8% of these expressed concern, compared to only 17.6% with few worries. As before, the most popular objection was based purely on principle (23.5%). These individuals found such data sharing to be invasive. If they wish to limit collection, they could choose to change their permissions.

“I value my online privacy”, (#F).

Behaviour. Through analysing the two incidents, we calculated a *app score* of 4.2/5. This implies that participants are strongly concerned about their smartwatch

data. To assess empirical behaviour, we explored the pretest logs. Throughout this 18-day period, not a single participant had restricted their permissions. In contrast, these settings had been loosened by 90% of the users. This would reduce the constraints on convenient functionality [335]. Furthermore, two individuals installed additional apps with sensitive permissions⁵. This resulted in a behaviour *permission score* of 50.4%. Since training has not been received, this presents a worrying baseline. Although users expressed concerns, they readily shared their data.

RQ1. For this chapter’s first research question, we explored whether smartwatch users protect their data. Screen locks were quite common, with six participants enabling this feature. However, the other four had never tried a password. GPS was used by all individuals throughout the pretest period. This implies that location tracking might have been possible. In the case of permissions, behaviour was most concerning. While the settings were never tightened, they were frequently loosened. Based on these results, smartwatch users rarely take action to protect their data. Through our privacy game, we hope to encourage protective behaviour.

RQ2. Despite the lack of protection, participants reported privacy concerns. When considering stranger access, individuals expressed strong objections. They also disliked the notion of their apps being used. Although some enabled screen locks, those without responded similarly. Opinions of tracking were more varied, with concerns being moderate. However, when sharing was considered, participants were greatly opposed. Since GPS was constantly enabled, location details might be vulnerable. This disparity is supported by the location Paradox found by Menfors and Fernstedt [263]. Finally, our respondents were concerned by app access. They also rejected sharing between companies. Unfortunately, permissions became increasingly lax. Based on these results, concerns and behaviour seem misaligned. Therefore, the Paradox appears to be present. Our smartwatch game seeks to mitigate this issue.

8.5 Gameplay Period Findings

Overview. While participants were completing the pretest questionnaire, we installed games on their smartwatches. For the next 16 days, all users participated in the gameplay period. They played the games for five days, followed by a week-long pause. They then resumed interaction for four days, with this approach known to refine mental models [257]. Therefore, the apps were only used for nine of the days.

⁵As outlined in the Research Questions, ‘sensitive permissions’ denote those granting access to precise location or text message contents.

During this period, we monitored participants' usage of the games. This enabled us to analyse both performance and compliance. At the end of the period, the logs were downloaded from the watches. Participants also completed a brief questionnaire to evaluate their experience.

Smartwatch game usage

Performance. We first analysed how our individuals played the games. Excluding outliers, there was a mean duration of 6.2 minutes. In terms of number of plays, the average was 22.0. Since the apps should be used three times per day on nine 'active days', the game was played less frequently than specified. If the privacy game still encourages protection, it may be particularly persuasive.

In general, the users performed very well. All but two of the ten unlocked the highest difficulty, suggesting the challenges were frequently completed. However, both these individuals were part of the treatment group. We expected the privacy tasks to be challenging, and this appeared to be the case. Privacy settings often lack usability [59], particularly on smartwatches [183]. While this may be unintentional, vendors have an incentive to encourage data collection [334].

High scores were formidable, with a mean average of 993.1. The maximum total was 1500, but the researchers themselves could not get above 1100. Therefore, it was even more impressive that three participants passed this mark. This might suggest that our games were usable, indicating an improvement from the prototype. Although significance was impeded by our small sample, group means appeared to differ. While the treatment average was 885.4, the control group got 941.0. This again might imply that privacy settings are harder to use.

Challenges. By analysing the logs, we identified which features were found most challenging. The privacy challenge completion rates are illustrated in Figure 8.4.

In the generic game, many tasks were completed without fail. These included adjusting the alarm volume and enabling vibration. Such features should be simple to locate, since they clearly relate to the 'Sounds and notifications' menu. However, the completion rate was varied in the privacy game. Some tasks, such as disabling GPS ($\bar{x} = 97.8$) or checking permissions ($\bar{x} = 98.9$), were completed with ease. Again, these features were in descriptive menus: 'Connectivity' and 'Apps', respectively. However, screen locks were found to be challenging, with a pattern average of 84.0%. PINs were considered particularly complex, receiving only 53.1%. In our posttest interviews, participants found these features hardest to demonstrate. A screen lock

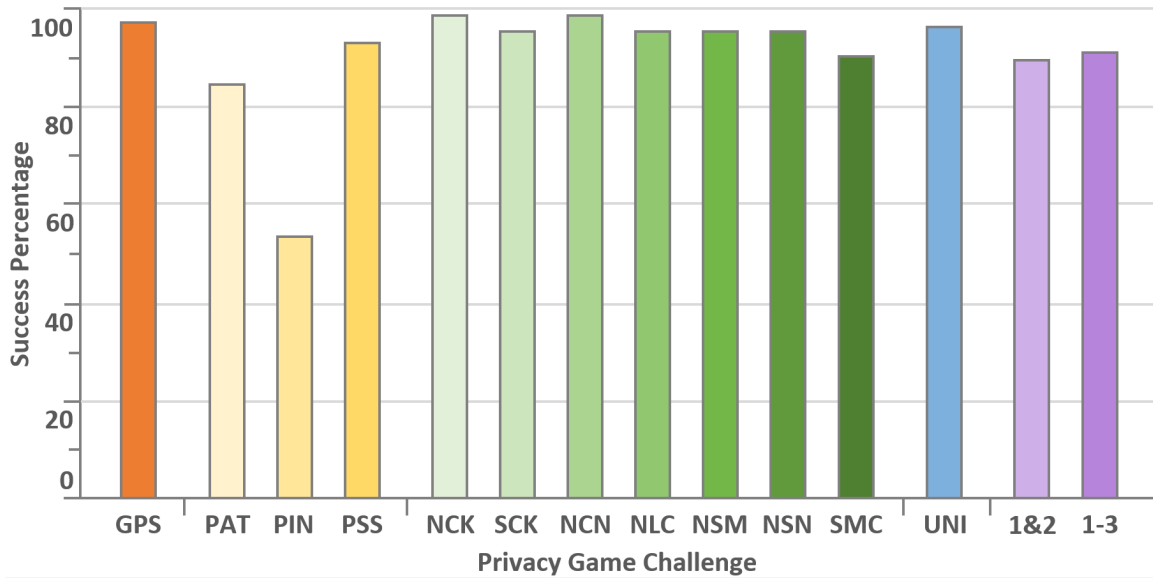


Figure 8.4: Privacy Game Challenge Completion Rates.

GPS; Pattern (PAT); PIN; Password (PSS); Permission Check (NCK); System App Permission Check (SCK); Revoke Contacts (NCN); Revoke Location (NLC); Revoke SMS (NSM); Revoke Sensors (NSN); Revoke System Microphone (SMC); Uninstall App (UNI); GPS and Pattern (1&2); GPS, PIN and Revoke Location (1-3).

might not intuitively belong within the ‘Personalisation’ menu. Again, this suggests that complex design can impede protective behaviour [59].

Gameplay evaluations

Overview. While the logs provided performance data, our questionnaire evaluated the user experience. The form was completed at the end of the gameplay period, while interaction was fresh in the memory. Through participant feedback, we can inform refinements for future apps. We can also assess whether our games were appreciated. Since they only differed in theme, experience differences might owe to privacy.

Evaluations. We began by asking whether the games were enjoyable. Since intrinsic motivation can be influential [331], we wished our apps to be appreciated. Across our sample, 60% expressed that they enjoyed the games. The games were also considered to be usable. 70% agreed with the statement, with none in opposition. Since the prototype faced interface issues, it appears these have been reduced.

Since we wished to inform participants, we hoped the apps were considered educational. Fortunately, all respondents agreed with the statement. However, differences existed between our groups. Whereas 80% of control participants were in strong agreement, this was matched by 20% of the treatment group. Surprisingly, this sug-

gests that the generic game was considered more educational. This opinion might have emerged for two reasons. Firstly, since the privacy app only concentrated on the three protective features, its content was narrower. Secondly, as evidenced in Question 5, privacy tasks were found more challenging. Whereas generic tasks were deemed simple (100%), nobody thought the same of treatment challenges. This contributed to a significant difference in perceived difficulty ($U = 0$, $p = 0.008$, $d = 2.928$). Finally, we asked users whether they learned anything from the game. Since 90% of our sample responded positively, it seems that both apps were educational.

Opinions. Our next questions solicited free-text responses. They sought to assess retention, inform refinements and explore veracity. We began by asking participants for their highest gameplay score. When reporting these figures, exaggerations might be encouraged by social desirability bias [138]. However, as shown above, we logged the actual figures. Fortunately, 90% named their score within a 10% range. Since accurate totals were reported, we have greater confidence in our participants' responses.

We moved on to solicit what individuals liked most and least. We also asked these users to suggest refinements. Across both groups, 17 positive comments were given. Usability was most praised, with this contributing to 41.2% of responses. The ease of interaction was particularly appreciated (23.5%), suggesting our game improved from the prototype. When discussing dislikes, 14 comments were provided. The most-frequent complaint was that the games were repetitive (28.6%). To refine knowledge on a topic, repetition is a standard approach [144]. Furthermore, we wanted participants to practise using the settings. In future implementations, greater randomness might make the repetition less apparent. 19 comments were provided as suggestions. 57.9% were in favour of extending the game, with 26.3% requesting additional challenges. This might suggest that users saw feasibility in our approach. Thus far, the permissions challenges focused on installed applications. In future versions, we could also consider the act of downloading apps. Permissions should be viewed in context, with some applications deemed less deserving of data [335]. By highlighting the risk of these programs, we might further encourage protective behaviour.

Finally, we assessed the participants' retention. We achieved this by requesting information about the games. We chose to ask for the house colours, since this landmark was common on all maps. Across both groups, the correct shades were named on 75% of occasions. It appears that our participants were retentive. In our treatment group, this might contribute to privacy education.

8.6 Posttest Findings

At the end of the 18-day posttest phase, participants each attended a final meeting. The smartwatches were returned and concern questionnaires were completed. Finally, we explored behavioural rationale through semi-structured interviews.

Participant opinions

Our posttest questionnaires explored both general opinions and privacy concerns. By comparing views against pretest responses, we can assess whether perceptions have changed. As before, the Opinions section could be subdivided into five themes: *attention*, *satisfaction*, *awareness*, *understanding* and *risk perception*. Since the treatment group have played our privacy game, we expect their knowledge to have increased.

Attention. Once again, attention was assessed through an Instructional Manipulation Check [286]. Since all participants gave the correct reply, they appeared to invest sufficient effort [286].

Satisfaction. We then explored whether the smartwatches were appreciated. Participants still found their devices useful, with 7/10 agreeing with this statement. Evaluations were slightly lower than in the pretest, and that might be because initial excitement has decreased. We then asked participants whether they used a range of functionality. Only 40% were in agreement, while 30% disagreed. At the pretest stage (Day 18), we thought users might have lacked time for exploration. If tools are not investigated after 8 weeks, many might be considered unnecessary.

While a product might be beneficial, this does not imply it has great importance. Therefore, we asked whether watch removal would cause inconvenience. Only two participants agreed with this statement, compared to 60% in opposition. Smartwatches might offer convenience, but they appear to be far from a necessity. If devices are not valued highly, perhaps their data will be deemed unimportant. Further research could explore whether this contributes to a lack of protection.

Awareness. We next explored the topic of privacy awareness. As in pretest, we asked users whether personal data could be read by apps. Both groups appeared to improve, with everybody agreeing with the statement (100%). Although treatment participants expressed stronger agreement, watch experience might have increased general awareness. If individuals now recognise the threat, they could reduce their risk by using protection.

Understanding. Since users rated their own knowledge, we also explore their self-efficacy. While 80% expressed confidence, most of these came from the treatment

group. Indeed, the only participants responding neutrally used the generic app. This might suggest that the privacy game can increase understanding.

Risk perception. We first asked whether a smartwatch might be lost or stolen. Since these devices are small and valuable, theft is possible [40]. In the pretest, 3/10 thought this risk existed. However, it had now increased to 50%, with only 20% in disagreement. If users perceive this threat, they can limit access with a screen lock.

To further explore the matter, we asked whether an app might threaten smartwatch data. Applications routinely access details [334], and permissions tend to be lax [59]. Whereas all the treatment group were in agreement (100%), 40% of the others disagreed. For those playing the privacy game, their mean increased from 3.0/5 to 4.8/5. Therefore, that app might have successfully highlighted risk.

Privacy concerns

Introduction. We now move on to assess the posttest concerns. Since individuals had now used their watches for 52 days, their views should have developed. They had also played one of two educational games. Since the privacy version sought to teach protection, concerns might be moderated by new knowledge. As before, these concerns were gauged through reactions to three pairs of contextualised questions: *stranger access*, *location tracking* and *app access*.

Stranger access. We began by requesting reactions to unauthorised data reading. Participants appeared concerned in general, with 80% expressing objections. However, the intensity differed between the two groups. While all the treatment group expressed concerns, 40% of the others were less convinced. The former group had played the privacy game, and hence might have learned of the risk. However, both groups had a mean of 4.2/5, since most controls were ‘Very Concerned’. This suggests that unauthorised access is still rejected. If users wish to decrease their risk, they could enable a screen lock.

Across both groups, we received 17 comments. 70.6% of these described the potential damage, suggesting that concerns were strong. For treatment individuals, the access to personal data was most troubling (40%). They also opposed the security risk that these details could pose (20%). In contrast, several control participants doubted their sensitivity (28.6%). However, if they knew their risk exposure, perhaps they would perceive a threat. Quotes, along with participant ID and group, are below:

“As it ... has data such as calendar, contacts and so on, it would not be the ideal scenario”, (#B, Treatment).

“Not sure they’d get much out of it”, (#I, Control).

Both groups continued to oppose unauthorised app use. 80% expressed concern at the issue, with distribution of responses being identical. This was greater than the 70% in pretest, suggesting the risk might have gained salience. The treatment group had a mean of 4.1/5, as did the control group. Across the two questions, the concern *stranger scores* were both 4.15/5. However, rationale differed when considering qualitative responses. 20 justifications were given, with 80% of them concerning damage. Our treatment participants named specific issues, such as impersonation (33.3%) or identity theft (22.2%). The control group were more general, and two expressed contingent concerns. Since our privacy game highlighted risk exposure, users might have gained knowledge of threats.

“Identity theft is my worst fear”, (#A, Treatment).

“If he uses my fitness app, then it is not a big deal”, (#G, Control).

Location tracking. As before, this issue was assessed through two incidents. Firstly, we solicited reactions to location tracking. At the pretest stage, only 50% were concerned by this issue. This now increased to 60%, since another treatment participant expressed their worries. No individuals were indifferent, resulting in a 3.6/5 mean for both groups. Although responses are not as strong as for some incidents, tracking still appears to be a concern.

Of the 18 qualitative justifications, 55.6% expressed unease. Treatment concerns were now less contingent, with some opposing the incident on principle (20%). Others were less worried, believing that the monitoring was for benign reasons (20%). Control participants feared the leak risk (25%) but cared less if it was optional (25%). With respondents expressing greater concern, we hope this contributes to GPS disabling.

“Because they could target me and I believe your location should remain private”, (#E, Treatment).

“There still runs a risk that there might be data leakage”, (#H, Control).

We also analysed reactions to location sharing. In the pretest, 80% expressed concern at this question. However, opinions had changed greatly by Day 52. The control group were still opposed, with all respondents being ‘Very Concerned’. In contrast, only two treatment participants acted in the same manner (40%). Indeed,

their concern appeared to decrease as the study progressed ($\bar{x} = 3.6/5$). Control users rejected the incident, represented by a maximal score ($\bar{x} = 5/5$). This resulted in concern *location scores* of 3.6/5 and 4.0/5, respectively.

To investigate this contrast, we analysed the qualitative responses. 15 justifications were provided, with 60% highlighting the issues. In the control group, individuals feared a security risk (33.3%). One participant was also worried because they felt they were uninformed (11.1%). While only three responses were in mitigation, all three came from the treatment group. Some thought tracking was benign (16.7%) or a modern reality (16.7%). Another claimed to now use GPS rarely and therefore face a lesser risk (16.7%). Their report was true, and this action might have been encouraged by the privacy game. Since their concerns were moderated by their behaviour, these two factors might have realigned.

“They can only do so if I have my location turned on, and as I only use this feature occasionally it wouldn’t bother me too much”, (#B, Treatment).

“In bad hands may be used against my interest”, (#J, Control).

App access. Finally, we requested reactions to an app accessing personal data. Respondents appeared to be concerned, since 70% expressed objections. However, views again differed between our groups. All the control group were worried, with 40% giving strong responses. In contrast, 40% of treatment participants gave neutral replies. While this does not imply indifference, concerns did appear slightly moderated. This contributed to a treatment mean of 3.8/5 and a control score of 4.4/5.

Of the 15 justifications, 60% identified potential dangers. The control group were worried about data selling (16.7%) and the risk of leakage (16.7%). Although treatment participants objected to the principle (22.2%), they were the only to offer mitigative views. It was expressed that data could be collected through alternative methods (11.1%). While true, permissions are an available means to limit access.

“Companies already have means of getting so not concerned”, (#D, Treatment).

“Wouldn’t trust the companies to not sell the information to other companies”, (#F, Control).

For the final incident, we assessed reactions to app data sharing. This provoked concern in the pretest, where 80% expressed their opposition. However, as in the previous question, treatment-group worries appeared to moderate. 80% of the control

participants were concerned at this scenario. This resulted in a posttest mean of 4.2/5. Treatment respondents were less worried, with 20% expressing indifference. Their mean ($\bar{x} = 3.6/5$) fell from the pretest, suggesting views became slightly moderated. If protection has increased in frequency, behaviour and concerns might be realigning. Our groups received concern *app scores* of 4.3/5 and 3.7/5, respectively.

Of the 16 qualitative justifications, 60% highlighted the potential risks. In the control group, targeted advertising was the main concern (37.5%). Since this is a common issue, it might have been known to the users. For treatment participants, reactions were contingent on other factors. Their concerns were nuanced, based on whether data was sensitive (16.7%) or aggregated (16.7%). As the game sought to provide education, perhaps users gained a detailed understanding.

“I wouldn’t mind if ... it was information that wasn’t specific”, (#B, Treatment).

“However, it may lead to annoying advertisements potentially”, (#H, Control).

Summary. In general, both groups continued to express concerns. When considering unauthorised access, opposition was strong. Individuals disliked the threats from tracking, impersonation and identity theft. If participants wish to reduce the risk, they could enable a screen lock.

When discussing location tracking, both groups had concerns. However, responses differed when considering sharing. Control users feared the security risk, with one complaining they felt uninformed. The treatment group were less worried, particularly since a participant had disabled GPS. In this manner, concerns appeared to adjust based on protective behaviour.

Reactions also differed when considering the actions of apps. The control group were worried about data selling and the risk of leakage. However, treatment-group worries appeared to moderate. They expressed nuanced views, considering whether access was sensitive or beneficial. Through their gameplay knowledge, they might have made informed decisions. If protective behaviour has also increased in frequency, the Paradox might be reduced.

Protective behaviour

Overview. By comparing pretest and posttest activity, we can analyse whether our games might have impact. We can also explore this chapter’s third research question: *Can the smartwatch game encourage privacy-protective behaviour?*. Since

our treatment group used the privacy version, we expect their protection to increase. If control behaviour is static, our questionnaires should have introduced little bias.

A per-participant comparison can be found below in Table 8.5. Increases in protective behaviour are highlighted in green. Reductions are in red, with white cells denoting little change. Please note that while we seek a reduction in GPS usage and permissions acceptance, we desire an increase in screen locking. Although treatment-group protection was not perfect, behaviour change is rarely absolute. Indeed, we were pleased that tools continued to be used at the end of the posttest period.

Table 8.5: Participant Changes in Protective Behaviour
Green: Increase in Protective Behaviour, Orange: Decrease in Protective Behaviour

#	Screen Lock	GPS	Permissions
Treatment	A	-	+1.1%
	B	+99.9%	-22%
	C	-	+21.5%
	D	+13.7%	-100%
	E	+100%	-
Control	F	-	-
	G	-	-
	H	-	+0.1%
	I	-	+1.1%
	J	-	+0.4%

Screen locks. In the treatment group, the pretest *lock score* was 57.3%. After gameplay, this rose to 100%, indicating a 42.7% increase. While our sample size impeded significance, we received a ‘medium’ effect size ($d = 0.733$) [333]. 38 participants would be required for 0.8 power, and this was infeasible for our study. The effect size implies that a considerable influence was present. This supports the increase in protection found in our Chapter 7 prototype. Control-group behaviour failed to change, suggesting our methodology did not encourage alterations.

In the pretest period, screen lock usage was far from standard. Six participants never used the feature, while one used it intermittently. However, the treatment-group change was dramatic. Within 10 minutes of playing the game, *B* and *E* had enabled protection. This suggests that the application was persuasive. While *D* had used the feature before, now their behaviour became consistent. After one day of gameplay,

all the group had the lock enabled. The protection was not removed, even three weeks after app interaction. These participants were concerned about unauthorised access. Their usage of protection increased after playing the game. In this manner, the Paradox might have been mitigated.

As outlined above, the control group failed to change. Their *lock score* was 20% in pretest and 20% in posttest. *J* enabled a password at the start of the study. Their behaviour was consistent and the lock was used throughout the process. However, none of the other participants tried the feature across the 52 days. This presents a worrying baseline for untrained users. Although they feared unauthorised access, they did little to prevent it. This suggests their disparity might still be present.

GPS. In the pretest period, GPS usage was constant. Throughout the 18-day phase, not a single person disabled the service. Since location questions triggered concern, a disparity might have been present. However, following the gameplay period, behaviour greatly differed. In the treatment group, the *GPS score* decreased by 40% from 100% to 60%. Although our sample size impeded significance, the ‘very large’ effect size was promising ($d = 1.461$) [333]. In contrast, no control participants disabled the service. This suggests our privacy game might have been persuasive.

Treatment-group behaviour changed dramatically. All participants used GPS during the pretest period. However, within two hours of playing the game, *B* disabled the service. Participant *D* also turned off GPS during the gameplay phase. Encouragingly, neither of these individuals re-enabled the functionality. This implies, even three weeks after interaction, that lessons were still retained. Concerns appeared moderated in posttest, partially in response to increased protection. As the two factors slightly realigned, the Paradox may have been mitigated.

In the control group, behaviour failed to change. Throughout the pretest period, GPS was constantly used (100%). This continued in posttest (100%), with the service never disabled in the 52 days. This was despite this group expressing more concern. Such a result has three implications. Firstly, the disparity appears prevalent for these users. Secondly, if they are not shown features, then settings might not be used. Finally, our methodology does not appear to have biased behaviour.

App permissions. Throughout the pretest period, not a single permission was restricted. Indeed, these settings were loosened by 90% of users. Despite the lack of protection, individuals expressed strong concern. These concerns remained strong for the control group. Treatment participants appeared less worried, balancing privacy against convenience. In terms of behaviour, both groups were reasonably static.

Across the sample, 80% acted mostly as they did before. Group comparisons are expanded below.

For treatment participants, there were some alterations. *B* appeared to respond to their concerns, decreasing their acceptance by 21.6%. As permissions were not used before, they might have learned from the game. However, *C* continued to install apps and loosen these settings. As will be discussed in the interviews, this individual claimed to balance privacy and functionality. Therefore, rather than making accidental mistakes, their behaviour was informed. The *permission score* barely differed, from 48.9% in pretest to 49.5 in posttest. It should be noted that if *C* was omitted, we might have reduced acceptance by 4.7%. Although the game did not strengthen permissions, it might have informed participants. With treatment-group concerns becoming nuanced, this is a possibility.

While there was some variation in the treatment group, control participants acted consistently. None had restricted a permission in the pretest period. Since some apps were not opened, this contributed to a *permission score* of 51.8%. This behaviour continued throughout the rest of the study (52.1%). Therefore, permissions were not restricted once across the two-month duration. This might be since individuals seldom expect superfluous data to be collected [343]. Users repeatedly reported strong privacy concerns. However, this seemed to have little influence on actions. As before, our findings have several implications. Firstly, a concern-behaviour disparity appears to exist within this group. Secondly, protection might be neglected if users are not prompted. Finally, since the control game did not encourage changes, the privacy version appears fairly evaluated.

RQ3. For this chapter's third research question, we explored whether privacy protection could be encouraged. Screen lock usage increased by 42.7% in our treatment group. Furthermore, those who enabled the feature did not disable it. The use of GPS fell sharply (-40%), with the service not being re-enabled. Although permissions acceptance was static, one individual restricted their settings. Behaviour differed greatly in the control group. Screen lock usage did not vary from pretest to posttest. GPS was enabled for all 52 days, with this setting never adjusted. Finally, permissions acceptance continued to be high throughout the study. We accept that treatment-group protection did not become perfect. However, interventions rarely have 100% efficacy, particularly when behaviour is assessed over several months. Since protective actions increased in the treatment group, our privacy game appears persuasive.

Table 8.6: Longitudinal Privacy-Protective Behaviour.
 GPS usage, Screen Lock usage and Permissions acceptance.
GPS scores: < 49% green and > 50% orange.
 (Screen) *lock scores*: < 49% orange and > 50% green.
Permission scores: < 34% green, 34% - 66% yellow and > 66% orange.

#		PRETEST PERIOD	GAME	PAUSE	GAME	POSTTEST PERIOD
TREATMENT GROUP	A	G				
		S				
		P				
	B	G				
		S				
		P				
	C	G				
		S				
		P				
	D	G				
		S				
		P				
	E	G				
		S				
		P				
CONTROL GROUP	F	G				
		S				
		P				
	G	G				
		S				
		P				
	H	G				
		S				
		P				
	I	G				
		S				
		P				
J	G					
	S					
	P					

Longitudinal behaviour

Introduction. To illustrate the changes in behaviour, we produced the graphic in Table 8.6. Columns represent time and are split into the pretest, gameplay and posttest periods. Rows denote participants and their privacy-protective behaviour. The top half showcases the five individuals in the treatment group. The bottom half presents the control users. Again, it should be noted that assignment was randomised before the study commenced. For each person, the usage of the features is shown. Since we desired low *GPS scores*, they are denoted in green. For the *lock scores*, green represents the frequent use of protection. Finally, since *permission scores* are less binary in nature, we selected three colours. In this case, green represents a low percentage. By comparing our groups, we can illustrate changes in behaviour.

Treatment. As can be seen in the top half, protection was rare in the pretest period. Although some participants used the features, their usage was inconsistent.

This is illustrated by the prevalence of orange shading. Once the gameplay phase begins, the chart becomes predominantly green and yellow. This remains throughout this period and the posttest phase. Even weeks after gameplay, protection did not appear to reduce. This suggests the longer-term feasibility of these applications.

Control. In the bottom half, the shading is predominantly consistent. Although two participants used some protection, the others neglected privacy features. This is illustrated by the prevalence of orange. When the gameplay phase begins, the behaviour fails to change. Indeed, actions appeared finalised from Day 2 of the study. Therefore, it appears that the generic game had no influence on protective behaviour. This implies that it served as an appropriate control.

RQ4. For this chapter's fourth research question, we explored whether the Paradox could be mitigated over the medium term. This is also the thesis' sixth sub-question and our central research question. To address this query, we must consider concerns and behaviour throughout the study.

At the end of the pretest period, participants reported their initial concerns. Across both groups, users feared access by a stranger. Although tracking faced moderate concerns, individuals objected to location sharing. Finally, participants were strongly opposed to the app access incidents. However, as we have outlined, protection was rare. Based on these results, behaviour and concerns appeared misaligned.

To assess whether the disparity was mitigated, we consider posttest responses. For concerns, we analyse the responses to the final questionnaire. For behaviour, we also consider actions on the final day. Since this was almost three weeks after gameplay, the privacy app might have lost salience. However, if concerns and behaviour are better-aligned, we appear to have mitigated the Paradox over the medium term.

We first consider the control group. Their concerns differed little as the study progressed. *Stranger access* triggered strong opposition, whether regarding data or app use. While *location tracking* was less concerning, location sharing was strongly rejected. Finally, this group was opposed to both the *app access* incidents. However, these claims did not lead to protective behaviour. Only one individual used a screen lock in pretest, and this continued. GPS usage and permissions acceptance also failed to change. Based on these results, both concerns and behaviour remained stable. Since they appear far from aligned, the Paradox is present.

In the treatment group, concerns became less extreme. When considering *stranger access*, the participants continued to show opposition. However, *location tracking* failed to provoke a strong reaction. Whereas users rejected sharing before gameplay, they now appeared less concerned. This was partially due to GPS disabling, which

would have reduced the risk [314]. Concerns also decreased when considering *app access*. Treatment participants offered nuanced justifications, mentioning aggregation and sensitivity. Across these scenarios, concerns appeared to slightly moderate.

This was matched by an increase in the usage of protection. While two participants used a screen lock at Day 18 (pretest), all five had it enabled at Day 52 (posttest). This could reduce the risk of unauthorised access. GPS usage also changed as the study progressed. Although the service was constant in pretest, two individuals later switched it off. Through this, location tracking should have been limited. Finally, we considered the degree of permissions acceptance. The percentage increased for one participant and decreased for another. It remained stable for the other users. In general, protective behaviour gained frequency in the treatment group. Concerns were moderated, often based on protection, so these two factors became better aligned. Since this was true 18 days after gameplay, we argue that we mitigated the Paradox over the medium term. **This addresses RQ4, the sixth subquestion and our thesis' central research question.**

The adjustments in concern and behaviour are illustrated in Figures 8.5 (treatment) and 8.6 (control). Both the *GPS score* and *permission score* are inverted, so that increases in the y-axis indicate protection. The external bars represent the pretest results, whereas the internal bars denote the posttest metrics. As shown, concerns in the treatment group tended to moderate. Although the *stranger score* increased, it was matched by a high usage of screen locks. In terms of behaviour, the group took greater action to protect themselves. Therefore, the two Paradox components appeared to realign.

Our control participants (Figure 8.6) were static throughout pretest and posttest. Although their *stranger score* slightly decreased, concerns and behaviour were generally consistent. This suggests that the privacy game was the persuasive factor. Medium-term Paradox mitigation has never been previously achieved, regardless of the environment. These findings also promote the feasibility of privacy games. With behaviour and concern better-aligned, user data should be at lesser risk. We now move on to explore rationale, in pursuit of RQ5.

8.7 Participant Rationale

Overview. While logs can outline behaviour, they do not provide rationale. For a rich analysis of perceptions, we conducted semi-structured posttest interviews.

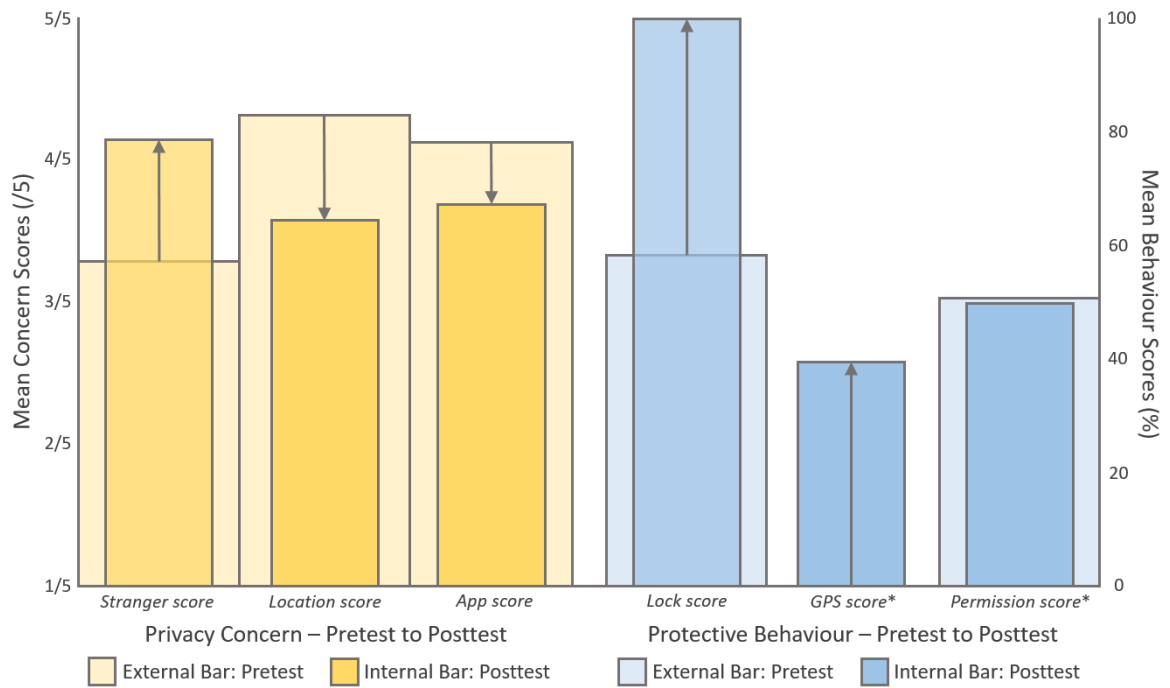


Figure 8.5: Privacy Paradox Changes: Treatment Group
*Scores inverted to represent GPS disabling and permissions denial.




Figure 8.6: Privacy Paradox Changes: Control Group
*Scores inverted to represent GPS disabling and permissions denial.

The interview was divided into four sections: *General*, *Awareness/knowledge*, *Protection Motivation Theory* and *Privacy Paradox*. The first two will be discussed in a

conventional fashion. The final two will be addressed through the participant profiles, which sought to characterise each individual. For brevity, we discuss the responses with central relevance to our research.

General responses

Motivation to participate. Although they were compensated, we hoped participants were motivated by smartwatch interest. Therefore, we asked them why they chose to take part. Of the 22 comments, 81.8% referred to trialling the devices. Indeed, all 10 were keen on investigating the watches. Only two mentioned the compensation, while two enjoyed academic research. Since users appeared interested, this should increase our external validity.

“I was about to buy a new one [smartwatch] so that was the perfect moment”,
(#J, Control).

Monitoring. We next asked users whether the background app might have affected their behaviour. If individuals were distracted by its presence, our findings would lose validity. Fortunately, no users believed the app had any effect. While this does not ensure external validity, it increases the reliability of our findings.

“I just used it as I would normally”, (#A, Treatment).

Learning. The interviews were conducted 18 days after the gameplay period. To assess retention, we asked users if they learned anything. Privacy was unprompted at this point, so references should suggest an organic interest. The topic was highlighted by 60% of the treatment group. Since these individuals also praised the game, it might have been informative. Controls were similarly influenced, with 60% mentioning their app. Regardless of the topic, it appears our games were educational.

“I think there was a couple of privacy settings that, through the game, I picked up that hadn't really hit me before”, (#B, Treatment).

Awareness/knowledge responses

General awareness. To assess awareness, we asked participants how they thought their data could be accessed. This explores the risk element of the concept. The treatment group appeared informed, with all users explaining the practices of apps. They also highlighted the risk of user accounts (11.1%) and fraudulent applications

(5.6%). In contrast, only two controls knew of app practices. This mirrors the low awareness found in our Chapter 6 interviews. Since users had watches for 52 days, the threat might lack salience. With awareness greater in the treatment group, this might have been provided by the game.

“Through some app that you allow them to track your location or your data”, (#E, Treatment).

General protection. To explore knowledge, users were asked what they would do to protect their privacy. We did not describe particular threats, since we wished to receive organic responses. In our treatment group, all participants mentioned a protective feature. 40% named screen locks, 60% cited permissions and 60% would disable GPS. Since all 11 comments cited a correct tool, our game appears educational. Even if individuals choose not to act, at least they should make informed decisions. The control group possess less knowledge, with only 40% able to name a single setting. One suggested reducing their notifications (6.7%), while others were unsure (13.3%). Since their approaches offer little protection, they might be placed at risk.

“I’d probably start by going through the list of apps again and seeing what permissions were useless”, (#A, Treatment).

Stranger access. We then explored responses to the contextualised questions, beginning with unauthorised access. All the treatment group mentioned screen locks, whether PIN (60%) or pattern (40%). This suggests they have learned how to respond to risks. Therefore, the disparity between concern and behaviour should be reduced. The screen lock rate was only 60% in the control group, with others suggesting credential changes (8.3%) or proximity disconnection (8.3%). Since 60% were unsure, these users appear less prepared.

“Again I haven’t looked at that”, (#I, Control).

Location tracking. We then explored behavioural responses to location monitoring. All the treatment group mentioned the importance of disabling GPS (100%). A further 80% cited permissions, as these provide additional protection. It appears as if the privacy game was considered educational. In contrast, controls appeared unprepared. Nobody referenced watch GPS, though some mentioned smartphone location (28.6%). As 60% were unsure, this further signifies the contrast in knowledge.

“Just switch off the GPS, easy. Settings and location and then you switch it off. That’s it”, (#C, Treatment).

App access. Finally, we solicited responses to our final scenario. The treatment group were again prepared, since all five would restrict permissions (100%). 60% also mentioned app deletion, implying users now know how to mitigate risks. Again, this should reduce any disparity between concern and behaviour. Unfortunately, no controls cited permissions, even after 52 days of smartwatch use. 80% were unsure, with only one suggesting app deletion. When education is not provided, individuals might lack the skills to protect themselves.

“You could revoke permissions for your contacts or your personal data, etcetera”, (#E, Treatment).

Screen lock demo. Individuals might know of a technique but be unable to use it. Therefore, we asked participants to demonstrate the three protective features. By talking aloud, we ascertained both their route and their confidence. Participants were first asked to enable the screen lock. For this setting, individuals had some difficulty. 70% followed a direct path, while 30% navigated through an indirect route. This appeared due to the placement of the feature in the ‘Personalisation’ menu. Our game did appear to encourage confidence, with 80% of treatment participants displaying certainty. In contrast, all the control group expressed confusion.

“I don’t see anything in these things so it’s not in Sound, I guess it’s Personalisation or Accessibility”, (#G, Control).

GPS demo. The treatment group easily found the GPS, with all five navigating directly (100%). They also expressed certainty (100%), likely because they are accustomed to the feature. Therefore, if they did want to restrict monitoring, they could. Despite control participants using a watch for 52 days, only 40% directly located the service. 80% were uncertain, with half of them claiming a lack of experience. If individuals are not accustomed to a feature, they might fail to adjust it.

“Disable GPS you said? Go down to Connectivity, Location, off. Done”, (#A, Treatment).

Permissions demo. When testing permissions, 80% of our treatment group were successful. The other individual was confident, but experienced technical difficulties. Results were more varied for controls. Although 80% navigated correctly, 60% lacked confidence in the route. Permissions are often opaque [35], and this might impede navigation. One user also claimed ignorance of this feature. Based on the group contrast, our game appears to support protective behaviour. Since users can respond to their concerns, the Paradox should be less prevalent.

“I go into Settings, go down to Apps, and System Apps. Google Play for example. Go down to Permissions, and then it shows you which ones”, (#E, Treatment).

Participant profiles: Treatment

Introduction. Thus far, we have studied rationale at a high level. To understand why individuals do (or do not) protect their privacy, we now analyse responses through PMT [328]. We questioned participants on all six components, and also explored justifications for the Paradox. Through considering the influential factors, we can gain novel insights into smartwatch behaviour. We begin by discussing the treatment group, with views potentially impacted by the privacy game.

Profiles. When comparing groups, coding proportions can be useful. However, perceptions are specific to particular individuals. Therefore, to discuss rationale, we constructed participant profiles. These rich characterisations were based on the replies to PMT and Paradox questions. Since their accuracy was crucial, they were each validated by their respective respondent. Through these profiles, we provide insight into user perceptions. For brevity, we discuss the participants with distinct and archetypal views. The remaining profiles can be found in Appendix C.

Profile A. They understood how to protect their privacy, but sometimes refrained to allow functionality. Data access was considered alarming when it was unexpected. When it provided features, it was judged to be less troubling. In terms of vulnerability, they were aware of app practices. However, the participant did appreciate the benefits of fitness applications. They were able to change their settings effectively, and believed they were quick. They also had some confidence in feature protection, but believed nothing was totally safe. Knowledge was no issue; rather, they chose to exchange certain pieces of data. Indeed, they believed their watch would have little use if everything was restricted.

“Yes, I think that’s the main reason we let them have our data, because they give us some sort of use”.

Profile C. This was a curious case - *C* had strong concerns but used an array of apps. What was most important to them was awareness and control. Access was considered severe if unconsented, and still serious if the rewards were not commensurate. They were fully aware of the vulnerabilities, recognising companies profit through data mining. The participant appreciated app functionality and hence restricted those providing little benefit. They were confident when it came to the efficacy of settings and their own ability to use them. They also believed the required effort was low, but still more than most would invest. The Paradox was justified on the distractions of watch functionality. It can lead people to forget about data collection. They believed that control was key, and wished that permissions were more fine-grained.

“I’m aware of when people, companies, whatever can access my data. I know that sometimes it’s the price that is to be paid for something”.

Profile D. *D* increased their protection, but accepted data collection as a modern reality. Due to this fact, they did not think watch access was highly severe. They also received benefits from sharing, believing that apps are useful when personalised. Through playing the game, the participant became knowledgeable and skilled in protection. They also had good awareness of data collection practices. The participant was confident in both the features and their ability to use them. They believed that these tools were quick, but that people need awareness first. When justifying the Paradox, they described how humans only act when threatened. While they recognised sharing was a social norm, they learned the skills to protect their data.

“I think human beings as a rule are generally quite passive, and faced with something that doesn’t have an active impact on us, we put it off”.

Group overview. In the interviews, treatment participants demonstrated that they could use their settings. Despite this fact, protection was not absolute after the gameplay period. This did not appear due to a lack of awareness. PMT discussions revealed that individuals consciously traded some data for benefits. For example, they might grant permissions to trusted apps. But they were aware of risks and could prevent access if necessary. This was highlighted by the use of locks and restriction of GPS. This is in great contrast to those who lack awareness and protective skills.

Our game sought to emphasise the risks, and participants acknowledged the privacy threat. It also aimed to support protective responses, and efficacy appeared strong. Based on this, we believe the game succeeded in its purpose. As participants became informed, objections decreased and protection became stronger. This suggests that concerns and behaviour became better-aligned.

Participant profiles: Control

Profiles. We now discuss the behavioural rationale of our control group. We found that data was frequently considered to be non-sensitive. This might be true, and could justify a lack of protection. However, users appeared to lack awareness, and this can impede the perception of risk [350]. Furthermore, individuals tend to underestimate privacy threats [220]. To explore the true sensitivity of data, we analysed the behaviour logs. These detailed the apps installed and their current permissions. We then compared this information against the user justifications. If personal details are possessed, this might demonstrate a further lack of awareness. We discuss two profiles of interest, with the others found in Appendix C.

Profile G. Since *G* did not consider their data to be sensitive, they neglected their protection. They still deemed access to be serious, as they disliked the practices of tech companies. They were fully aware of collection and received benefits from calendar reminders. However, in non-smartwatch environments, they avoided functionality to protect their data. After they saw the features, they were confident in their efficacy and their use. The tools were seen as low effort, but inconvenient if functionality is desired. They might have investigated protective features if they deemed their data to be sensitive. The participant justified the Paradox on people caring in principle rather than in practice.

“I am concerned, but I do minimum things to say that I’m concerned. And I’m trying to do something about it, but I still don’t want to rid myself of the benefits”.

G claimed to neglect the features since their data was innocuous. However, their watch could reveal some information. Phone contacts were available through a Google application. With passwords not enabled, a stranger might learn of *G*’s friends. These people could then be stalked or targeted. An installed fitness app, *Spartan Body Weight Workouts*, might not store sensitive data. But it would give strangers an insight into the participant’s personality. We then considered the permissions of default apps. Through the *Fit* tool, Huawei could monitor location and heartrate. With GPS never disabled, the collection could occur frequently. While *G* doubted sensitivity, their watch could reveal some details.

Profile I. This user differed in their logic, neglecting protection due to a lack of salience. Consent was their key concept, with data access expected if an app was installed. They had knowledge of collection practices and received benefits from fitness applications. However, they expressed that they easily overlooked protection.

These settings were seen to be low effort, but private defaults would be even easier. The participant believed that even though features were helpful, they lacked total control. They also thought they might use settings if prompted by the watch. The Paradox was justified on people assuming their data is private. If the topic is made salient, users might be encouraged to protect themselves.

“I think people kinda assume that the default is going to be in your interest, rather than in the company’s interest”.

Their watch possessed many pieces of personal data. *Facebook Messenger* can house phone contacts and private conversations. *WearCalendar* was also installed, which stores reminders and appointments. While *Spotify* might not reveal sensitive data, it provides insight into a personality. Since a screen lock was never enabled, data could be accessed if the watch was mislaid. Company access was also possible, as GPS was constantly on. Location could be monitored through either *Accuweather*, *Fit* or *Google Maps*. The permissions for all these apps were lax. Therefore, watch details were quite accessible.

Group overview. Whereas the treatment group made informed decisions, the control group appeared to lack privacy knowledge. Most doubted the sensitivity of their data and did not perceive a risk. As evidenced from the interviews, many also had little understanding of protection. This suggests a worrying baseline for untrained smartwatch users. In contrast, treatment participants recognised the value of information, and controlled the details they exchanged. The privacy game did not result in a total rejection of data-accessing apps. Instead, it has provided information so that concerns and behaviour could be better aligned.

Smartwatch behaviour drivers

Finally, following the construction of the profiles, we reflected on the rationale. While certain considerations were common, others were infrequent. By analysing these factors, we can gain greater insight into smartwatch behaviour.

PMT factors. Within the treatment group, participants generally possessed good self-efficacy. 55.6% of their comments cited confidence, compared to only 25% for the control users. However, after learning of the techniques, our sample deemed response costs to be low. Indeed, 10/10 participants believed the features were simple to change. We also explored response efficacy, with 80% trusting that settings provide protection. Therefore, the privacy settings do not appear the primary issue.

The threat components seemed to be more influential. Firstly, most users had a balanced view of severity. Of the 31 opinions relating to this topic, 48.4% were contingent on other factors. If data access was consented and rewarded, many were satisfied. This helps to explain why features were often not used. Secondly, the control group failed to perceive risk. As outlined earlier, only 2/5 were aware of app practices. Furthermore, unlike treatment users, they deemed their data to be innocuous. Finally, participants received rewards from smartwatch functionality. 80% indicated that they exchanged some information for convenience. Since settings can impede useful apps [34], permissions might be accepted.

Rationale. Based on our participant discussions, three issues appeared to influence decisions. While other factors might be important, these considerations were prevalent. They comprised: *sensitivity*, *salience* and *convenience*.

When individuals knew that their data was valuable, they considered protection. For example, Participant *J* enabled a screen lock due to their array of personal apps. This further emphasises the importance of risk perception, as highlighted in our Chapter 6 interviews. However, since some deemed details to be innocuous, settings were underexplored. To further encourage protection, the value of data should be highlighted. This could be achieved through illustrating company processing practices. If users recognise their data's worth, they should make informed decisions.

Since privacy often lacked salience, some participants forgot the concept. They might have been preoccupied with other functionality or immediate tasks. This was also suggested through our Chapter 6 interviews. However, if users felt at risk or noticed consequences, features regained their relevance. These consequences were highlighted through the game's interactive challenges. It might be difficult to keep privacy salient on smartwatches. The topic is often a secondary goal [188] and settings are hidden within complex menus. If watch faces could display current configurations, that might improve the situation.

Crucially, users tended to weigh convenience against protection. Smartwatches are obtained to provide functionality, and settings can restrict these benefits. As found in the Chapter 5 survey, IoT purchases are often encouraged by exciting features. But individuals were informed, they could make a delicate balance. For example, Participant *C* would disclose some data, but only if they appreciated the advantages. This indicates that privacy calculus has some influence over decisions. However, the calculus will only be accurate if users are informed [350]. Indeed, when individuals lack knowledge, they tend to overestimate the benefits [157]. Therefore, it will remain important to provide individuals with privacy awareness. To assist in imparting

knowledge, and potentially mitigating the Paradox, we suggest the development of further educational games.

8.8 Implications

Paradox. Through this final study, we addressed our central research question. We sought to explore whether the Paradox could be mitigated over the medium term. This was achieved through a two-month longitudinal study. During this period, we empirically analysed smartwatch behaviour. We also solicited participant opinions at several points. Our treatment group received a privacy game, refined based on the prototype findings. Control participants used another application [344], as we sought to reduce confounding variables. In the former case, concerns appeared moderated as the study progressed. This might have been in response to improved knowledge or behaviour changes. Protective settings were also used more frequently, reducing the disparity. In contrast, responses in the control group were static. Throughout the 52-day period, these participants never disabled GPS or restricted a permission. This inaction suggests that protection was not due to our methodology. Based on our findings, the privacy game appeared to mitigate the Paradox. As behaviour remained strong over several weeks, we addressed our central research question.

Rationale. To understand the rationale behind the Paradox, we conducted posttest interviews. The treatment group displayed sound privacy awareness, and were able to demonstrate protection. Control participants expressed little understanding and struggled to use the settings. Three considerations appeared to influence behaviour: sensitivity, salience and convenience. Without privacy awareness, individuals tended to undervalue their data. Protection was also overlooked, frequently because functionality was more compelling. Importantly, convenience encouraged the loosening of settings. However, when users were aware of the risks, these decisions were made in an informed manner. Participants would guard other data and had the knowledge to protect themselves. We consider this to be a success of the thesis. While we mitigated the Paradox, protection will never be absolute. But when individuals possess knowledge and self-efficacy, they are empowered to take accountability.

Chapter 9

Conclusions

9.1 Introduction

The purpose of this thesis was to investigate the Privacy Paradox within a smart-watch environment. Through the development of educational games, we appeared to reduce its prevalence. In this final chapter, we conclude our research. We begin by summarising our work and demonstrating it addresses our research subquestions. Since all studies have limitations, we then critically assess our findings. After this critique, we discuss our contributions to research. Penultimately, further work is outlined based on these contributions. Finally, we conclude with a high-level overview.

9.2 Summary

As with most projects, it is crucial that our work addressed the research subquestions. Therefore, we summarise our studies in terms of addressing these points. Our central query was as follows: *Can the Privacy Paradox be mitigated in the context of smartwatches?* Due to the growth of this novel environment, we argue this was an important issue to tackle. Although some individuals might balance risks against reward, calculations tend to be biased by a lack of awareness [157]. Users can place themselves at inadvertent risk [194] and then express regret at their actions [385]. Through our five novel studies, we addressed our central research question.

Can the Privacy Paradox be confirmed in the UK? While the Paradox has been studied in several nations, perceptions can differ by culture [18]. Therefore, to support later analyses, we confirmed its presence in the UK. This was primarily achieved through a street survey across four geographic locations (Chapter 4). Claims were strong, with over 90% of respondents expressing concern for privacy. Over 80% also reported that they acted privately with their data. However, online protection

was used inconsistently. Furthermore, participants readily disclosed their details, even when this disclosure was optional. There appeared to be little relationship between concerns and behaviour. Since the Paradox was corroborated in Chapters 6, 7 and 8, we believe it was confirmed in the UK.

How do perceptions differ between IoT and other devices? Since the IoT is often criticised in the media [131, 154, 273], we thought consumers might be apprehensive. A ‘fear of the unknown’ could also cause unease, increasing privacy concern [221]. If protective behaviour is constrained, the Paradox might be prevalent. To explore these issues, we compared perceptions across a wide range of technologies (Chapter 5). Three of these were from the consumer IoT, while three were older computers. Through an online survey, participants both indicated concerns and evaluated other factors. IoT products were deemed to be more troubling, with wearables provoking the most unease. They were also lacking in terms of usability and familiarity. This could constrain the usage of protective features [158, 230, 317]. Furthermore, privacy appeared rarely factored into purchasing decisions. Based on these considerations, we asserted that the IoT may be prone to the Paradox.

How does the prevalence of the Privacy Paradox differ between IoT and other devices? Chapter 5 suggested that the IoT provokes privacy concerns. Behaviour might also be constrained, but this was not explored in depth. For a rich analysis, we conducted interviews with 40 product owners (Chapter 6). 20 possessed IoT devices, while 20 owned less-novel computers. Through semi-structured discussions, we assessed reactions to potential risks. We also explored protective behaviour and the rationale behind it. Concerns appeared high in both groups. However, IoT protection was far less frequent, contributing to Paradox prevalence. Since this disparity was also confirmed in Chapter 7, we are confident that the IoT is susceptible.

Which factors contribute to the Privacy Paradox in the IoT? In our semi-structured interviews (Chapter 6), we did not solely assess behaviour. We also extracted the justifications behind these actions. To mitigate the Paradox, we must understand the factors that influence the issue. Fortunately, we explored the matter through a wide range of questions. Based on the full discussions, a lack (low degree) of awareness was the most-cited IoT reason. Awareness also appeared to affect disclosure behaviour in the Chapter 4 surveys. Many individuals either did not recognise the risk or the required protection. Other IoT factors included short-term necessity and poor risk perception. To mitigate the disparity, we sought to provide education and practice through games. In Chapters 7 and 8, the importance of awareness and risk was also found. Hence, we have confidence that these factors are influential.

Can we mitigate the short-term prevalence of the Privacy Paradox on smartwatches? Since wearables appeared particularly prone to the Paradox, we then scoped to smartwatches. To provide awareness and practice, we sought to develop educational games. However, since this context is novel and challenging, a prototype appeared prudent. Therefore, we began by assessing whether mitigation was possible over the short term. We developed the first application to consider smartwatch privacy (Chapter 7). This was evaluated and refined through an online study with 504 smartwatch owners. Most participants appeared concerned about their privacy. The control group did not play the game and expressed static behaviour. In contrast, treatment participants increased their protection after gameplay. As a result, the Paradox appeared mitigated. This was the first work to reduce the issue over several weeks. With feasibility demonstrated in the short term, we then targeted a longer period.

Can we mitigate the medium-term prevalence of the Privacy Paradox on smartwatches? In Chapter 8, we sought to address our central research question. To evaluate the Paradox over the medium-term, we conducted a two-month longitudinal study. This was the first such analysis to target this phenomenon. Refined through Chapter 7 feedback, we developed the first privacy game for (Android) Wear OS. It sought to highlight awareness and risk exposure through dynamically-customised challenges. Five treatment participants received this app, while our five controls played a generic version. Their protective behaviour was empirically monitored throughout the 52-day period. In the control group, concerns were high and protection was rare. After the gameplay phase, the Paradox appeared to decrease for treatment participants. Behaviour did not relapse for the final three weeks, and awareness appeared sharp in posttest interviews. Therefore, we argue that the issue was mitigated over the medium term. Through this, and the preceding studies, we addressed our central research question.

9.3 Critique

Overview. Through our aforementioned work, we believe we addressed our research questions. However, no studies are without their limitations. Therefore, we now critically reflect on our conducted work. We argue that despite areas for refinement, we provide novel and valuable insights.

Can the Privacy Paradox be confirmed in the UK? In Chapter 4, we confirmed the existence of the Paradox in the UK. To achieve this, we studied the

correlation between concerns and behaviour. Since no significant relationship was found, we concluded that the phenomenon was present. However, a lack of significance does not guarantee that an effect is absent. Significance can be limited by sample size, and this is likely in a 112-person study. Null hypothesis testing is contentious [277], and we accept the aforementioned points. To minimise the risk of a type I error, we also analysed the correlation coefficient. This is akin to an effect size and is not influenced by the size of a sample [159]. In three (H3, H4 and H5) of the four correlation hypotheses, r_s was less than ± 0.1 . This denotes a very weak relationship [366], which should not be expected between concern and behaviour. For H6, the coefficient was also only -0.146, suggesting little connection. We deemed this evidence of the Paradox, and our findings were corroborated in later chapters.

We studied disclosure behaviour as a proxy for privacy protection. However, if revealed demographics were false, participants did not act in a lax manner. For ethical reasons, we did not verify the accuracy of data. Nevertheless, we believe that fabrication should be infrequent. It was easier to mark ‘Prefer not to say’ than to complete the fields. Furthermore, since respondents gave informed consent, the interaction may appear official. While verification could have provided a ground truth, the ethical issues are challenging. As a partial solution, we could solicit truthfulness in follow-up forms, as done by Metzger [266].

As previously discussed, most participants disclosed their demographic data. However, if details are deemed innocuous, behaviour might not differ from concerns. We accept that not all individuals will value their marital status, Twitter username and employer name. Nevertheless, these details can reveal personal information. Twitter usernames are unique, with most accounts being publicly-accessible [255]. Employment is a professional matter, and a company name could aid identification. In future work, we could request more-sensitive details. But this may deter participation from privacy-conscious individuals [164].

Finally, respondents might have trusted the researcher since he represented the University of Oxford. Through authority bias [267], this might have led to an increased willingness to disclose demographics. Although this is possible, we felt it inappropriate to disguise our affiliation. Furthermore, participants were informed verbally and textually that fields were optional. Finally, while authority might be commanded by a professor, this is less likely for a PhD student. To assess any potential influence, we could analyse disclosure with and without the affiliation.

Despite these limitations, 68.8% reported acting privately while revealing two details. Participants also failed to use consistent protection, despite claiming concern.

Even if some disclosed due to the above factors, the contrast remains strong. We confirmed Paradox existence and this was corroborated in Chapters 6, 7 and 8.

How do perceptions differ between IoT and other devices? In Chapter 5, we found that the IoT provokes concern. Since it was deemed to lack usability and familiarity, we believed it might constrain protection [158, 230, 247, 317]. However, since there is no dichotomy between ‘IoT’ and ‘non-IoT’, our analyses were challenged. While we agree that a spectrum exists, a baseline group was required for smart device comparisons. Marginal products, such as mobile phones, were also excluded from our study. Furthermore, even if we ignore ‘non-IoT’ analyses, the IoT provokes concern. In a refined work, we would use interviews to define the category boundaries.

Since many products are not considered, some might deny that the categories represent the IoT. If true, this would undermine Paradox exploration in this environment. We accept that several devices, such as connected vehicles, are not studied. However, for a practical and feasible analysis, we compared six products. While we could have analysed a wider range, this might have contributed to response fatigue [88]. Furthermore, we scoped to consumer devices as they should be better-known. Finally, categories were selected through a robust process. With a larger sample size, it might be possible to consider more devices. Each participant could then rate a subset, reducing the risk of fatigue.

To assess concerns, we asked individuals whether products respected their privacy. A negative response might not guarantee that a person is concerned. They might accept the risks and value other facets of the device. Although this is true, we wished to avoid ‘concern’ evaluations. The term might have primed our topic [65], so this question appeared a decent proxy. Furthermore, if a participant believes products show little respect, they likely express some concern. Fortunately, the presence of concerns was corroborated in Chapters 6, 7 and 8. In these cases, the matter was explored in a grounded manner through contextualised questions.

Although the evaluations supported our hypotheses, some IoT products were owned by few participants. Therefore, ratings might have been uninformed. Device popularity certainly varied, with only 12% possessing home automation systems. However, since we explored high-level perceptions at this stage, ownership was not a requirement. Furthermore, if users do fear unfamiliar technology, this is relevant to the novel IoT. While we considered surveying those with all six products, this would have constrained our sample. Therefore, we explored the views of owners in the semi-structured interviews. In a refined survey, we could include a field for respondents to indicate their confidence.

Even if we ignore ‘non-IoT’ comparisons, the IoT appeared to pose a risk. Rather than evaluating a single product, participants rated wearables, smart appliances and home automation systems. In all three cases, they were deemed to lack respect for privacy. Furthermore, they were rated poorly in terms of usability and familiarity. It was impractical to consider all IoT devices, and our categories covered consumer technologies. As the Paradox was found in Chapters 6, 7 and 8, we have confidence in our findings.

How does the prevalence of the Privacy Paradox differ between IoT and other devices? Our Chapter 6 interviews suggested that the Paradox is more prevalent in the IoT. However, as behaviour was self-reported, it might have been fabricated. We accept that participants could exaggerate their protection. However, since only 22.5% reported reading policies, we deem this unlikely. If behaviour was exaggerated, this places a minimum bound on Paradox prevalence. Furthermore, by disguising the interview theme, demand characteristics should have been mitigated [287]. Empirical data is indeed preferable, and hence it was collected in Chapter 8.

Since concerns and behaviour were qualitative, we feared comparisons would be excessively-subjective. However, our assessment might have been impeded by the quantitative translation. For example, a person might not be ‘slightly concerned’ if they provide both positive and negative responses. In hindsight, we accept that translations should be more robust. We could have asked for numeric replies, followed by qualitative justifications. Despite these issues, we believe our findings are still sound. The analysis was consistent for all participants, and protection was certainly less frequent in the IoT. Interviewees also verified the appropriateness of their themes, increasing our confidence.

Category groups had different sizes, ranging from laptops ($n = 15$) to desktops ($n = 1$). IoT evaluations might have amounted to a wearable analysis, since 13/20 devices belonged to this group. We would have preferred equal sizes for comparisons. However, as demonstrated in Chapter 5, some technologies are less popular than others. This constrained recruitment for these rich in-person interviews. Furthermore, to consider the IoT holistically, it is appropriate that certain devices are more prevalent. In a refined study, we could conduct video interviews with equal groups.

When studying IoT devices, we assumed that privacy concerns and behaviour came from a single individual. This is often true, with the person configuring the device being the same that uses it. This would appear particularly likely in the case of smartwatches, since the product is attached to its owner’s wrist. However, this assumption is not accurate in all cases. When children use technology, their parents

might adjust the settings instead. Similarly, the elderly might receive assistance from younger relatives. In smart home environments, multiple individuals might interact with the same interface. We therefore accept that our work does not address privacy behaviour in more social or communal settings. Multi-user research would have challenged our Paradox analyses, since concerns and actions might belong to different parties. Furthermore, for ease of ethical procedures, we did not seek to interact with children. Future research should explore how communal interactions affect the Paradox.

There are certainly limitations in the evaluation process. However, interview responses suggest the Paradox is prevalent in the IoT. While passwords were used frequently on older technologies, they are often overlooked on smart devices. Privacy settings are configured occasionally and policies are seldom read. Despite this, concerns were high in both groups. Since tools were not used in response to risks, the Paradox appeared prevalent. Chapters 7 and 8 also corroborate that the issue is common.

Which factors contribute to the Privacy Paradox in the IoT? In these interviews, we also explored the considerations which influence protective behaviour. By understanding these factors, we designed approaches to mitigate the Paradox. However, the justifications were those the participants thought relevant. They might not be true, and this could undermine our developments. Although we accept this, we do explore the rationale of IoT owners. Furthermore, rather than being directed by their suggestions, we considered prevalent themes. Finally, a lack of awareness remained common in Chapter 7 and 8 control groups. Since the Paradox decreased in treatment groups, this suggests the factor had some influence.

To explore the IoT, we conducted interviews with 20 smart device owners. Since they were recruited locally, the sample is unlikely to be representative of IoT users. For a greater analysis of rationale, we could have targeted a larger number and range of individuals. While this is true, few samples could truly represent the diversity of opinions. To extract the contributory factors, we opted for depth rather than breadth. This was achieved through our detailed discussions with device owners. We do not claim to have identified all the influential concepts. However, awareness was also a common issue when evaluating 504 smartwatch users (Chapter 7).

As previously mentioned, a lack of awareness appeared to be the most prevalent IoT issue. This was one reason why we developed educational games. However, the frequency of a factor does not imply its influence. Other concepts might be more

important but mentioned on fewer occasions. Furthermore, ignorance might be sometimes claimed to disguise apathy. We accept these points, and agree that frequency is one of several considerations. We also do not assert that privacy awareness is the only issue. Nevertheless, participants repeatedly used it to justify the Paradox. Its influence is supported by the literature [48,309], and education appeared a feasible solution. Based on our Chapter 7 and 8 results, we are confident it was influential.

We accept these limitations to our approach. However, we believe we extracted a range of contributory factors. Indeed, risk perception was also targeted in the smartwatch games. Our detailed interviews were conducted with those owning IoT devices. All questions were contextualised around these products, and the responses were coded through an inductive process. Respondent validation then increased the likelihood that our findings aligned with the data [56]. A lack of awareness was common, as supported by previous literature [49,247]. Since education appeared to mitigate the Paradox, we believe our evaluation was appropriate.

Can we mitigate the short-term prevalence of the Privacy Paradox on smartwatches? Through our Chapter 7 evaluation, we found that the prototype encouraged protective behaviour. This suggested that the Paradox could be mitigated over the short term. However, since actions were self-reported, they might have been fabricated. While we accept that this is possible, we took approaches to mitigate the risk. Firstly, to reduce demand characteristics [287], privacy was disguised in both the advertisement and the questionnaires. Secondly, participants were anonymous and compensated regardless of their responses. Finally, if protection was exaggerated, it should affect both groups. Since 34.5% still displayed the Paradox in posttest, it appears unlikely that behaviour was inflated. We agree that empirical data is preferable, and the longitudinal study corroborates our findings.

Since the Paradox decreased in the treatment group, this was attributed to the privacy game. However, since participants watched a brief video, this might have had some influence. The video indeed concerned the topic of privacy. But it was used to introduce the content of the game. Due to the drag/swipe interactions of our online prototype, we feared participants would have difficulty. While an in-person demonstration would be preferable, this was impractical in this remote study. Again, due to the success of Chapter 8, the efficacy of games is corroborated. In a revised design, we would exclude the video and include greater textual description.

Although protection seemed to increase, it was measured on a frequency scale. A person might act infrequently but still guard their data. For example, permissions may be adjusted rarely if few apps are installed. Furthermore, permissions are

relative to the requirements of each particular application [335]. However, for most settings, usage frequency has a relationship with degree of protection. If a lock is used consistently, unauthorised access should be limited. Furthermore, the more GPS is enabled, the greater the period that data might be read [180]. For a quantitative analysis, this scale appeared apt. It was preferable to binary questions, which might be reductive.

Since our prototype was evaluated remotely, there were limitations to the work. However, the Paradox did appear to be mitigated. Concerns were high in both groups, and remained high after the gameplay session. Control behaviour was static, with these participants not using the prototype. In contrast, the treatment group improved in terms of lock usage and permissions adjustment. With these actions reducing the concern risks, the disparity appeared to decrease. The longitudinal study corroborated that games can be persuasive.

Can we mitigate the medium-term prevalence of the Privacy Paradox on smartwatches? In our longitudinal study, our privacy game appeared to mitigate the Paradox. However, participants received a loaned smartwatch and were aware of device monitoring. Therefore, due to social desirability bias [138], they might have used more protection than normal. We accept this risk, and the Hawthorne Effect is a common bias [7]. Nevertheless, we doubt that behaviour was influenced. The monitoring app resided in the background and did not appear to disrupt the users. Since we also received approval for a range of data, we sought to disguise our purpose. We could have hidden our collection or studied the usage of owned devices. However, both approaches presented ethical issues. As a partial solution, the participants could have kept the watches. If they feel ownership, they should display genuine behaviour.

The control group were found to rarely protect their watches. But since participants took part in an academic study, they might have believed their data was safe. While this risk is accepted, it is likely common to much privacy/security research. Individuals were encouraged to use the devices as normal. Furthermore, 40% of participants enabled screen locks in pretest, and this is unlikely if protection is assumed. The information sheet, read by all users, detailed potential vendor access. However, we felt that behaviour would be biased if risks were clearly highlighted. As treatment users adopted protection, it is unlikely they believed data was automatically safe.

Screen lock usage and GPS usage were assessed in terms of frequency. In contrast, permissions were analysed based on the acceptance rate. Although we only considered sensitive permissions, this rate does not consider the context. For example, while it may be acceptable for maps to request location, a news app might be resisted

[335]. We accept this limitation and reserve a detailed exploration to future work. However, we did discuss the context of permissions within the posttest interviews. With thousands of apps available, it would be challenging to assess each application. Furthermore, Wear OS does not support analysis of when data is accessed. However, if we study open source watches, such as AsteroidOS¹, we might have fewer constraints.

We accept the limitations posed by our methodology. However, even if these issues had some influence, protection still increased in the treatment group. Concerns also slightly moderated as the study progressed. In contrast, control participants barely changed their responses or actions. This suggests that our methodology did not promote the alterations. Since the disparity was mitigated in this two-month study, we believe our research question was addressed.

9.4 Contributions

Through our work, we produced several contributions to research. In some cases, these directly aligned with the subquestions. In others, they emerged organically through our five studies. By highlighting consistent themes throughout the chapters, we now outline novel insights.

Awareness is low in IoT/smartwatch environments. Previous work has found that individuals tend to have low privacy awareness [49, 247]. When users do not recognise the risk, they might not guard their data. Through our research, we found that awareness is particularly lacking in the IoT. For example, it was a common theme within our semi-structured interviews (Chapter 6). 40% of those with smart devices were unsure who could access their data. Of greater concern, 35% claimed to have never considered the matter. When justifying the Paradox, a lack of awareness was the most-common IoT theme. The issue was also prevalent in smartwatch environments, as found in our prototype evaluation (Chapter 7). 48% of controls reported not understanding protection. A further 44% lacked confidence in usage, with over half blaming insufficient skills. The longitudinal study confirmed this low awareness (Chapter 8). After 52 days, 60% of controls had little idea how to protect their data. The same proportion had no knowledge of device access. It appears that understanding is particularly poor in these environments.

This is likely due to the novelty and unfamiliarity of the fields. Furthermore, since privacy can be a secondary goal [188], users might be preoccupied with other features. As the norm for connectivity grows, this preoccupation is unlikely to decrease [41].

¹<https://asteroidos.org/>

Without support and education, as found in our research, individuals might not protect themselves. This has implications for both research and government. If the IoT poses such confusion, tools should be developed to highlight privacy. ‘Nutrition labels’ have been effective in other environments [205], and tools could be used to simplify policies [241]. Companies have few incentives to raise awareness. While the GDPR seeks to curb invasive practices, we call on governments to take additional steps. For example, IoT products should be mandated to include accessible instructions. These would detail not only the privacy policy, but how to restrict settings. If users are better-supported, they can make informed decisions.

IoT data is deemed to be innocuous. As previously mentioned, the IoT could collect large quantities of data [300]. This collection can occur through physical sensors or virtual interactions. Due to the ubiquity of the environment, details might reveal sensitive information. However, our research suggests that IoT data is often deemed innocuous. These opinions first arose in the interviews, with 35% citing the fact when discussing deletion. It was also mentioned in response to the sharing (35%) and monitoring incidents (20%). Furthermore, of those neglecting passwords, 37.5% blamed a lack of sensitivity. In the prototype evaluation, we found smartwatch opinions were similar. Of the controls who doubted the threat, 45.7% believed data was innocuous. Few thought privacy was not worth it, but 30% of these cited the factor. Finally, when we interviewed controls in the longitudinal study, 60% doubted sensitivity. Therefore, IoT data appears to be assigned a low value.

Evaluations might be accurate, particularly if details are transient or abstract. However, since users tend to lack privacy awareness, they might not appreciate possible inferences [106]. Furthermore, as collection can be surreptitious [273], the stored data might be unknown. Our findings would be of interest to privacy research, particularly the PETRAS project². Their Value of Personal Data in IoT (VPD) theme seeks to explore perceptions of smart device metrics. Future analyses could explore how valuations vary across a spectrum of products. Low evaluations also have several further implications. Firstly, if details are deemed useless, protection might be neglected. Since this suggests an IoT privacy risk, research should seek to highlight sensitivity. Secondly, users might readily sacrifice their data if it is undervalued. This partially explains the popularity of free-but-invasive services [334]. If company practices could be illustrated, consumers might appreciate their details. Finally, when data is not valued, the device itself might be disregarded. This poses a risk to vendors, with our longitudinal users judging watches to be unnecessary. The IoT might

²<https://www.petrashub.org/>

lead to increased obsolescence, with products rapidly replaced. In this environment, privacy may become an afterthought.

Protection is rare in IoT/smartwatch environments. Even when privacy protection is available, it tends to be rarely used. This has been previously demonstrated on smartphones [22] and social networks [4]. However, tool usage was unclear in novel environments. Through our research, we demonstrate that IoT protection is particularly low. This has important implications, as smart devices can collect large quantities of data [300]. The interviews first illuminated our user behaviour. ‘Non-IoT’ owners tended to set passwords and adjust their settings. However, the frequencies were 15% and 45% in the IoT group. Protection was also rare in smartwatches, as highlighted in our prototype evaluation. Only 24% used their permissions often, despite frequent GPS usage by most participants. When analysing empirical data, privacy settings were even less accessed. Over 52 days, our controls never disabled their GPS or restricted a permission. This behaviour appears even more lax than that in other environments [22]. Despite settings being available, interviews suggested they were overlooked.

Although settings can be time-consuming, the response costs were not deemed to be high. More often, individuals were not aware of tool’s presence. Users might also be preoccupied with functionality, hence reducing the salience of privacy. It was claimed that protection would be used if data was sensitive. However, this appears optimistic since judgements are constrained [350]. If IoT settings are frequently avoided, this has implications for research and industry. As privacy lacks salience, academics should seek to highlight the topic. This might reduce preoccupation, and success has been found on social networks [189]. Since settings can limit functionality, research should also explore how to grant data while minimising sensitivity. For industry, we have demonstrated the importance of interface design. Screen locks were categorised within an unexpected menu (‘Personalisation’), and hence they were challenging to find. We suggest that privacy features are centralised within their own settings. Although the topic can be in tension with business models, recent controversies could make privacy vogue [156]. Temporarily, there might be a competitive advantage to offering a semblance of control.

Smartwatch behaviour is influenced by salience, sensitivity and convenience. Through our final two studies, we explored the behavioural rationale of smartwatch users. In both cases, several factors appeared to be influential. If privacy was not salient, individuals seemed less likely to use protection. This was demonstrated by the role of awareness in the studies. Furthermore, while Participant *I*

regarded tools as simple, they often forgot to use them. Sensitivity appeared to have a strong relationship with risk perception. When threats were doubted or privacy deemed unnecessary, innocuous data was the largest factor (Chapter 7). Individuals also claimed to be influenced by the presence of risk. Finally, since smartwatches are consumer devices, actions are often undertaken for convenience. This factor drove IoT purchases (Chapter 5) and was a common justification for lax behaviour (Chapter 6). Even when our Chapter 8 participants became informed, they still balanced privacy against functionality. While other considerations also have influence, the above factors seem to impact behaviour.

These findings have novelty, as smartwatch privacy has not been deconstructed to such a degree. They are also significant, since factors could be adjusted to influence behaviour. We have demonstrated that awareness can encourage device protection. If users know of risks or responses, they should be more likely to act. Saliency could be increased by privacy reminders, with apps highlighting settings on a periodic basis. If sensitivity does have an impact, this property might be modified in several ways. Vendors may wish to understate the value, seeking to encourage disclosure. However, as watches begin to house additional information, users might investigate protection. Finally, companies will be pleased that convenience drives interactions. If firms can develop exciting features, consumers remain likely to divulge. To resist this, researchers must design innovative approaches. MockDroid was successful in mobile environments, enabling functionality while masking sensitivity [54]. Smartwatch versions could support privacy alongside convenience.

Smartwatch games can encourage protective behaviour. In previous work, educational games were found to be persuasive [394]. This was true in a variety of security-based environments [199, 316, 341]. However, few privacy games were developed [363], with none targeting smartwatches. We first targeted this context through our prototype application. Although protection was far from perfect, the usage of screen locks and permissions increased. To corroborate this with empirical data, we developed the first smartwatch privacy game. While the control group failed to change, protection improved for treatment participants. Screen lock usage increased by 43% and GPS disabling rose by 40%. Based on qualitative responses, the app also enhanced awareness. Smartwatch games appear to be a persuasive technique.

This finding has many implications for research, industry and government. Privacy games are in need of further exploration, as only three have ever been developed (two of which are in this thesis). While smartwatches are a challenging environment, it is surprising that other platforms have not been targeted. We suggest that games

investigate different dimensions of privacy, such as confidentiality and anonymity. Partially informed by our longitudinal results, research should also extend educational techniques. Due to the small interface, Learning Science approaches should be adapted for smartwatches. App companies might also be interested in our findings. Games are growing in popularity on Wear OS and Apple Watches. If education is feasible, this presents another lucrative market. While our app targeted privacy, the theme could be applied to different topics. Google could include similar games to teach watch navigation. Finally, our behavioural findings have implications for government. With awareness campaigns often lacking efficacy [38], apps could serve as a complement. Based on game performance, officials could then monitor whether knowledge improves. We hope to promote a field of educational privacy apps.

The Paradox can be mitigated through awareness and education. As previously outlined, the Paradox has been explored in extensive detail [153, 214]. However, few works have sought to mitigate its prevalence. Jackson and Wang [194] addressed the issue in their 2018 study, but did not mitigate it over a period of time. Our research suggests that the disparity can be reduced through enhancing awareness. This was first indicated in our prototype evaluation, where the treatment group played a game. While control behaviour remained static, the Paradox reduced in the other group. Our longitudinal study corroborated these findings with empirical data. After users received education, their concerns appeared slightly moderated. Protective behaviour also increased and remained high for weeks after interaction. While not all individuals adjusted their settings, we supported informed decisions. We did not expect to eliminate the Paradox, but were pleased that the disparity decreased. This should reduce the likelihood that data is inadvertently placed at risk.

If our approach was successful, this has several implications for research. As asserted by Pöttsch [309] and Deuker [114], awareness does encourage protection. Therefore, to adjust the usage of services, greater information should be provided. Researchers could annotate other interfaces, helping users to make informed choices. We demonstrated the feasibility of this approach on smartwatches. With engagement constrained on these devices, education might be particularly successful in other contexts. Social networks [4] and smartphones [194] are prone to the Paradox, and they should face further analysis. Since our longitudinal study was enlightening, this methodology should be used to explore other factors. For example, we could vary functionality to assess its impact on protective behaviour. Our results suggest the Paradox can be mitigated. We hope this encourages further work to address the issue.

9.5 Future Work

Overview. Informed by our findings, there are several opportunities for future work. These can be divided into three themes: *Paradox exploration*, *smartwatch behaviour* and *educational games*.

Paradox exploration. In our longitudinal study, we analysed behaviour over a two-month period. Through our educational game, we appeared to mitigate the Paradox. However, it would be valuable to explore the matter over a longer term. If privacy loses salience, protective behaviour might decrease in frequency. By studying interactions over six months, we could gauge the true influence of the game. Furthermore, the Chapter 8 analysis was constrained to 10 participants. For a broader exploration, and to seek statistical significance, we would recruit a larger sample.

Our research suggests that Paradox prevalence can be reduced by increasing awareness. However, as demonstrated, some individuals will still exchange data for functionality. Since awareness is not a panacea, we should explore the behaviour of informed users. Participants could be taught the risks to privacy and the techniques for protection. Their knowledge could then be verified through an assessment. With awareness now removed as an issue, user behaviour should be fully informed. If actions still deviate from concerns, this might owe to trade-offs or cognitive biases.

Risk perception was cited as important in both the prototype evaluation and the longitudinal study. While our games highlighted risk exposure, they did not actually target user data. If individuals feel genuinely under threat, they might take protective action. Although the ethics are challenging, a threat could be simulated through several approaches. For example, by demonstrating inference techniques, users might learn the value of their data. This should reduce optimistic biases [29], and potentially lead to protection.

Smartwatch behaviour. Thus far, we have explored user actions at a high-level. In Chapters 7 and 8, the general usage of protection was considered. Since tools appear rarely used, a contextual analysis is now warranted. Certain apps might be used on weekdays but overlooked on weekends. Furthermore, permissions acceptance should vary based on the application and its data [335]. Previous work has explored negotiation rather than this ‘all-or-nothing’ approach, and this could be extended to smartwatches [34]. Wear OS products are limited in terms of the behaviour that can be monitored. By analysing ‘rooted’ devices, we could gain a greater understanding of how smartwatches are used.

Currently, many participants justified their lack of protection on innocuous data. Indeed, several claimed they would take action if their smartwatch was sensitive. However, it is unclear whether behaviour would actually adjust. To explore this, we would vary the information stored on a device. Without priming privacy, we would then analyse the usage of protection. If settings are consistently neglected, then sensitivity might have less influence than claimed.

Our games suggest that protection can be encouraged on smartwatches. With this confirmed, we can now explore techniques which require less engagement. Third-party watch faces are supported in many ecosystems. They could monitor current settings in a similar manner to our logging app. Colour-coded results could then be displayed, highlighting an individual's risk exposure. Through a 'gamified' approach, users may be persuaded to use protection.

Educational games. In this thesis, we explored the IoT before scoping to smartwatches. Since games appeared influential on the latter, perhaps they would support other smart devices. Smart TVs have been criticised for their eavesdropping and transmission of conversations [108]. As these products can host third-party apps, games could be developed to highlight privacy. Through a similar methodology to Chapter 8, we could assess their influence.

Our applications were implemented on Wear OS smartwatches. Although these devices are amenable to analysis, the wearable market is diverse. Apple Watches are highly popular and also support third-party apps. Therefore, similar games could be implemented in this environment. If protection can again be encouraged, smartwatch games might be a valuable resource.

9.6 Concluding Remarks

This work has explored the Privacy Paradox in the context of smartwatches. It provides this environment's first detailed analysis of protective behaviour. We first conducted 40 semi-structured interviews, providing a novel exploration of the IoT. Based on qualitative rationale, a lack of awareness (both risks and responses) appeared influential. While previous work suggested its importance [309], the Paradox had not been consistently addressed [194]. To mitigate the issue, we developed the first smartwatch privacy game. It was informed by a prototype evaluated by 504 smartwatch owners. Through a novel longitudinal study, we evaluated empirical behaviour over two months. The game appeared to encourage protection and address the Privacy Paradox. Since actions did not relapse, our approach seemed to be persuasive. This

should reduce the likelihood of data being inadvertently placed at risk. We encourage the development of further games to support informed privacy decisions.

Appendix A

Qualitative Coding Frames

It should be noted that top-level themes are sometimes defined based on discrete responses. For example, in the product perception surveys, participants indicated whether they owned a desktop or not. This was a discrete answer, with ‘Yes’, ‘No’ and ‘Unsure’ options available. The user then provided qualitative rationale, which was coded. To reduce the risk of justifications being misconstrued, the discrete responses formed the top-level themes. This was important, as syntactically-similar answers might support either confidence or insecurity. We highlight discrete top-level themes by appending an * to the table caption.

A.1 Chapter 5: Product Perception Surveys

Table A.1: Do you own a Desktop?*

Theme	Subtheme	Subsubtheme	Definition	Quote
Yes	Applications	Writing	Responses in this category indicate that a desktop was owned because it is helpful for writing.	<i>“Useful for word processing/editing”</i>
		Computing	Responses in this category indicate that a desktop was owned because it is helpful for programming tasks.	<i>“I am a software developer”</i>
		Gaming	Responses in this category indicate that a desktop was owned because it is helpful for gaming.	<i>“Use a PC for gaming”</i>
		Internet access	Responses in this category indicate that a desktop was owned because it is helpful for accessing Internet services.	<i>“Email, internet”</i>
		Research	Responses in this category indicate that a desktop was owned because it is helpful for research purposes.	<i>“It is essential for a researcher”</i>
		Image/video editing	Responses in this category indicate that a desktop was owned because it is helpful for editing/managing images and videos.	<i>“Best 4 photoshoppery”</i>
		Communication	Responses in this category indicate that a desktop was owned because it is helpful for communicating with others. This may be through specific platforms such as social media.	<i>“As a means of communication”</i>
		General entertainment	Responses in this category indicate that a desktop was owned because it supports leisure such as watching videos or listening to music.	<i>“Use for entertainment”</i>
		Storage	Responses in this category indicate that a desktop was owned because it is helpful for storing data.	<i>“In addition we store films on it”</i>
	Roles	Work	Responses in this category indicate that a desktop was owned to undertake work related to their job.	<i>“Required for work, business”</i>
		Studying	Responses in this category indicate that a desktop was owned to allow studying, possibly for university.	<i>“For studying primarily”</i>
		Daily standard usage	Responses in this category indicate that a desktop was owned because they are useful for daily tasks.	<i>“Helpful in daily life”</i>
	Usability	General usability	Responses in this category indicate that a desktop was owned because it is usable. However, rationale is not provided for why it is easy to use.	<i>“Easier to work on”</i>
		Large screen	Responses in this category indicate that a desktop was owned because it possessed a large usable monitor.	<i>“Having something with a bigger screen”</i>
		Keyboard	Responses in this category indicate that a desktop was owned because it possessed a usable keyboard.	<i>“Larger keyboard makes it easier to use”</i>
		Accessibility	Responses in this category indicate that a desktop was owned because it has a good level of accessibility.	<i>“Convenient for older people in my house”</i>

Do you own a Desktop?* *Continued*

	Advantages	Familiarity	Responses in this category indicate that a desktop was owned because it was familiar or remained from the past.	<i>"Familiarity due to my first ever tech"</i>
		Powerful	Responses in this category indicate that a desktop was owned because it has high processing power.	<i>"Heavy computing tasks"</i>
		Essential	Responses in this category indicate that a desktop was owned because computers are considered necessary.	<i>"It is required in today's modern world"</i>
		Cheap	Responses in this category indicate that a desktop was owned because the product is not too expensive.	<i>"At an affordable price"</i>
		Protect privacy	Responses in this category indicate that a desktop was owned because it can be configured to protect privacy.	<i>"Protect myself in term of privacy"</i>
	Unsure	Responses in this category indicate that a desktop was owned but the participant is unsure about the reason.	<i>"Unsure"</i>	
	Not applicable to response	Responses in this category are excessively concise or provide little rationale for their response. It might appear as if the reason supports the opposite answer.	<i>"It's there/why not"</i>	
Unsure	Security concerns	Responses in this category indicate that they are unsure whether they own a desktop, but they were reluctant due to security issues.	<i>"For security"</i>	
No	Prefer other devices	Prefer laptop	Responses in this category indicate that a desktop was not owned because the participant prefers using a laptop.	<i>"Superseded by laptops"</i>
		Prefer tablet	Responses in this category indicate that a desktop was not owned because the participant prefers using a tablet.	<i>"Work tablet"</i>
		Prefer smartphone	Responses in this category indicate that a desktop was not owned because the participant prefers using a smartphone.	<i>"Smartphone"</i>
	Disadvantages	Too large	Responses in this category indicate that a desktop was not owned because it was deemed to be too large and bulky, hence consuming excessive amounts of space.	<i>"It takes up too much space"</i>
		Not portable	Responses in this category indicate that a desktop was not owned because it lacked mobility and portability.	<i>"It's not portable enough"</i>
		Too expensive	Responses in this category indicate that a desktop was not owned because it was deemed too expensive.	<i>"Because it's expensive"</i>
		Too old	Responses in this category indicate that a desktop was not owned because it was deemed too old.	<i>"Too old"</i>
	No need	Responses in this category indicate that a desktop was not owned because the participant did not need such a device.	<i>"Nothing I do requires the use of a desktop computer"</i>	

Do you own a Desktop?* *Continued*

	Owned by workplace	Responses in this category indicate that a desktop was not owned by the participant, but instead provided by their workplace.	<i>“Owned by my work not me”</i>
--	--------------------	---	----------------------------------

Table A.2: Do you own a Laptop?*

Theme	Subtheme	Subsubtheme	Definition	Quote
Yes	Applications	Writing	Responses in this category indicate that a laptop was owned because it is helpful for writing.	<i>“I use it for word processing”</i>
		Computing	Responses in this category indicate that a laptop was owned because it is helpful for programming tasks.	<i>“I can program on it”</i>
		Internet access	Responses in this category indicate that a laptop was owned because it is helpful for accessing Internet services.	<i>“Useful for internet browsing”</i>
		Image/video editing	Responses in this category indicate that a laptop was owned because it is helpful for editing/managing images and videos.	<i>“Image and video processing”</i>
		Communication	Responses in this category indicate that a laptop was owned because it is helpful for communicating with others. This may be through specific platforms such as social media.	<i>“Communications”</i>
		Entertainment	Responses in this category indicate that a laptop was owned because it supports leisure such as watching videos or listening to music.	<i>“Watching TV, music”</i>
		Gaming	Responses in this category indicate that a laptop was owned because it is helpful for gaming.	<i>“Gaming”</i>
		Data syncing	Responses in this category indicate that a smart appliance was owned because it allows data to be synched across a range of different devices.	<i>“Like it syncs with other Apple devices I own”</i>
	Roles	Work	Responses in this category indicate that a laptop was owned to undertake work related to their job.	<i>“Too complete work based tasks”</i>
		Studying	Responses in this category indicate that a laptop was owned to allow studying, possibly for university.	<i>“Use it for assignments for university”</i>
	Usability	General usability	Responses in this category indicate that a laptop was owned because it is usable. However, rationale is not provided for why it is easy to use.	<i>“It’s more user friendly than a desktop”</i>
		Large screen	Responses in this category indicate that a laptop was owned because it possessed a large screen.	<i>“Big screen to watch programmes”</i>
		Keyboard	Responses in this category indicate that a laptop was owned because it possessed a usable keyboard.	<i>“Useful as has full sized keyboard”</i>

Do you own a Laptop?*

Continued

	Advantages		Familiarity	Responses in this category indicate that a laptop was owned because it was familiar or remained from the past.	<i>"Familiar"</i>
			Powerful	Responses in this category indicate that a laptop was owned because it has high processing power.	<i>"Sufficient processing power to run all the software I need"</i>
			Essential	Responses in this category indicate that a laptop was owned because computers are considered necessary.	<i>"One cannot function without"</i>
			Cheap	Responses in this category indicate that a laptop was owned because the product is not too expensive.	<i>"It's very practical and more cheap than other technologies"</i>
			Portable	Responses in this category indicate that a laptop was owned due to its mobility and portability.	<i>"Convenience in portability"</i>
			Small	Responses in this category indicate that a laptop was used due to its compact size.	<i>"Takes up less space at home"</i>
			Convenient	Responses in this category indicate that a laptop was used due to its convenient functions.	<i>"Laptop is convenient"</i>
			Protect privacy	Responses in this category indicate that a laptop was owned because it can be configured to protect privacy.	<i>"Protect myself in term of privacy"</i>
			Flexible	Responses in this category indicate that a laptop was owned because it can be used for a flexible range of tasks. It can also be customised to suit many purposes.	<i>"Flexible, customizable"</i>
			Received		Responses in this category indicate that a laptop was used because it was either received from another person or possessed by a family member.
		Not applicable to response		Responses in this category are excessively concise or provide little rationale for their response. It might appear as if the reason supports the opposite answer.	<i>"N/a"</i>
Unsure	Not true ownership		Responses in this category indicate that they are unsure whether they own a laptop, since it is unclear whether they own all the components within the device.	<i>"I bought my laptop, but technology is too broader term for me to truly own it"</i>	
No	Prefer other devices	Prefer desktop	Responses in this category indicate that a laptop was not owned because the participant prefers using a desktop.	<i>"Prefer desktop pc"</i>	
		Prefer tablet	Responses in this category indicate that a laptop was not owned because the participant prefers using a tablet.	<i>"Have a tablet with wireless keyboard instead"</i>	
	Too expensive		Responses in this category indicate that a laptop was not owned because it was deemed too expensive.	<i>"Useful but expensive"</i>	
	No need		Responses in this category indicate that a laptop was not owned because the participant did not need such a device.	<i>"No need to own one now"</i>	

Do you own a Laptop?* *Continued*

	Owned by workplace	Responses in this category indicate that a laptop was not owned by the participant, but instead provided by their workplace.	<i>“Use work laptop”</i>
	Not applicable to response	Responses in this category are excessively concise or provide little rationale for their response. It might appear as if the reason supports the opposite answer.	<i>“Not everything can be done on a smartphone”</i>

Table A.3: Do you own a Tablet?*

Theme	Subtheme	Subsubtheme	Definition	Quote
Yes	Applications	Reading	Responses in this category indicate that a tablet was owned because it is helpful for reading books and documents.	<i>“Kindle to read”</i>
		Entertainment	Responses in this category indicate that a tablet was owned because it supports leisure such as watching videos or listening to music.	<i>“Entertainment value”</i>
		Internet access	Responses in this category indicate that a tablet was owned because it is helpful for accessing Internet services.	<i>“Internet access everywhere”</i>
		Apps	Responses in this category indicate that a tablet was owned because it possesses useful and fun apps.	<i>“Range of apps”</i>
		Gaming	Responses in this category indicate that a tablet was owned because it supports gaming.	<i>“Plus I like the silly games”</i>
		Communication	Responses in this category indicate that a tablet was owned because it is helpful for communicating with others. This may be through specific platforms such as social media.	<i>“Connecting with friends/family”</i>
		Writing	Responses in this category indicate that a tablet was owned because it is helpful for writing.	<i>“From writing essays”</i>
	Roles	Work	Responses in this category indicate that a tablet was owned to undertake work related to their job.	<i>“Easy to use in work”</i>
		Studying	Responses in this category indicate that a tablet was owned to allow studying, possibly for university.	<i>“Studying”</i>
	Advantages	Powerful	Responses in this category indicate that a tablet was owned because it has high processing power.	<i>“Given that mine runs full Windows, more power-hungry functions as well”</i>
		Quick convenience	Responses in this category indicate that a tablet was used due to its quick and convenient functionality/productivity.	<i>“Because it is super convenient”</i>
		Portable	Responses in this category indicate that a tablet was owned due to its mobility and portability.	<i>“Access to technology anywhere you go”</i>

Do you own a Tablet?* *Continued*

		Readable screen	Responses in this category indicate that a tablet was owned because it possessed a readable screen.	<i>"Larger screen than smart-phone"</i>
		Usable	Responses in this category indicate that a tablet was owned because it is usable.	<i>"Very easy to use"</i>
		Small	Responses in this category indicate that a tablet was used due to its compact size.	<i>"It doesn't take up too much space"</i>
		Cheap	Responses in this category indicate that a tablet was owned because the product is not too expensive.	<i>"Cheap"</i>
	Needed for children		Responses in this category indicate that a tablet was owned since it provided support with their children.	<i>"To keep kids busy"</i>
	Received		Responses in this category indicate that a tablet was owned because it was either received from another person or possessed by a family member.	<i>"Husband bought me a tablet"</i>
	Despite privacy concerns		Responses in this category indicate that a tablet was owned but the participant still possessed privacy concerns.	<i>"But not 100% give user full control and customisation of privacy"</i>
	Unsure		Responses in this category indicate that a tablet was owned but the participant is unsure about the reason.	<i>"Unsure"</i>
	Not applicable to response		Responses in this category are excessively concise or provide little rationale for their response. It might appear as if the reason supports the opposite answer.	<i>"N/a"</i>
No	Prefer other device	Prefer smartphone	Responses in this category indicate that a tablet was not owned because the participant prefers using a smartphone.	<i>"I use my smart phone instead"</i>
		Prefer laptop	Responses in this category indicate that a tablet was not owned because the participant prefers using a laptop.	<i>"Laptops suffice for my life"</i>
		Prefer desktop	Responses in this category indicate that a tablet was not owned because the participant prefers using a desktop computer.	<i>"Prefer desktop"</i>
		Device not specified	Responses in this category indicate that a tablet was not owned because the participant prefers using other non-specified devices.	<i>"Use alternatives"</i>
	Disadvantages	Poor keyboard	Responses in this category indicate that a tablet was not owned because the keyboard was considered to be small or difficult to use.	<i>"Not big enough to be able to type on"</i>
		Too large	Responses in this category indicate that a tablet was not owned because it was deemed to be too large.	<i>"Too big to go in a pocket"</i>
		Too expensive	Responses in this category indicate that a tablet was not owned because it was deemed too expensive.	<i>"Haven't money"</i>
	Unrequired	No need	Responses in this category indicate that a tablet was not owned because the participant did not need such a device.	<i>"I have no need for it"</i>

Do you own a Tablet?* *Continued*

	Object to technology	Responses in this category indicate that a smart home device was not owned because the participant dislikes this type of technology.	<i>"I hate tablets"</i>
	Received	Responses in this category indicate that the participant did not consider themselves to own a tablet as they did not purchase it themselves.	<i>"Husband iPad"</i>
	Not applicable to response	Responses in this category are excessively concise or provide little rationale for their response. It might appear as if the reason supports the opposite answer.	<i>"I don't have one"</i>

Table A.4: Do you own a Wearable?*

Theme	Subtheme	Subsubtheme	Definition	Quote
Yes	Applications	Fitness tracking	Responses in this category indicate that a wearable device was owned because it supports heart rate monitoring and fitness tracking.	<i>"To track my health and fitness"</i>
		Timekeeping	Responses in this category indicate that a wearable device was owned because it allows the time to be monitored easily.	<i>"I like the way it displayed the time"</i>
		Appointments	Responses in this category indicate that a wearable device was owned because it allows calendars, appointments and reminders to be displayed.	<i>"Setting reminders"</i>
		Notifications	Responses in this category indicate that a wearable device was owned because it provides useful messages and notifications.	<i>"Android wear device: useful companion to phone notifications"</i>
		Portable connectivity	Responses in this category indicate that a wearable device was owned because it allows online functionality to be accessed in a portable manner.	<i>"Access to technology anywhere you go"</i>
		Remote control	Responses in this category indicate that a wearable device was owned because it allows products to be controlled remotely.	<i>"Remote control of some functions"</i>
	Advantages	Quick functionality	Responses in this category indicate that a wearable device was owned because it supports quick and convenient functionality.	<i>"Very useful for quick replies to messages"</i>
		Usable	Responses in this category indicate that a wearable device was owned because it is usable.	<i>"Easy of use"</i>
		Experimenting	Responses in this category indicate that a wearable device was owned because the participant is experimenting with different forms of technology.	<i>"To test how usable smart-watches are at the moment"</i>
		Received	Responses in this category indicate that a wearable device was owned because it was either received from another person or possessed by a family member.	<i>"Purchased for me as a gift"</i>

Do you own a Wearable?* *Continued*

	Despite privacy concerns		Responses in this category indicate that a wearable device was owned but the participant still had privacy concerns.	<i>"Privacy concerns around data generated, rather than devices themselves per se"</i>
	Unsure		Responses in this category indicate that a wearable device was owned but the participant is unsure about the reason.	<i>"Unsure"</i>
	Not applicable to response		Responses in this category are excessively concise or provide little rationale for their response. It might appear as if the reason supports the opposite answer.	<i>"Na"</i>
No	Prefer other device	Prefer smartphone	Responses in this category indicate that a wearable device was not owned because the participant prefers using a smartphone.	<i>"I already have a phone"</i>
		Prefer non-connected watch	Responses in this category indicate that a wearable device was not owned because the participant prefers using a normal non-connected timepiece.	<i>"I have a proper watch"</i>
	Unrequired	No need	Responses in this category indicate that a wearable device was not owned because the participant did not need such a device.	<i>"I don't need one"</i>
		Object to the technology	Responses in this category indicate that a wearable device was not owned because the participant dislikes this type of technology.	<i>"I think it's silly and not necessary. Just a naff consumerist gadget"</i>
	Costly	Too expensive	Responses in this category indicate that a wearable device was not owned because it was deemed too expensive.	<i>"Too expensive"</i>
		Little resale value	Responses in this category indicate that a wearable device was not owned because they were considered to have poor second-hand resale value.	<i>"And no second-hand market value"</i>
	Disadvantages	Poor usability	Responses in this category indicate that a wearable device was not owned because it was deemed to have poor usability.	<i>"No easy interface"</i>
		Underdeveloped	Responses in this category indicate that a smart appliance was not owned because it was deemed to not be functional enough.	<i>"It is not advanced enough to be worth buying"</i>
		Small screen	Responses in this category indicate that a wearable device was not owned because it was deemed to have an excessively small interface.	<i>"Too small of a screen"</i>
		Too bulky	Responses in this category indicate that a wearable device was not owned because it was deemed to be too large and bulky on the participant's wrist.	<i>"Looks a bit too bulky and masculine"</i>
Ugly		Responses in this category indicate that a wearable device was not owned because they are considered to look unattractive.	<i>"Because it is ugly"</i>	

Do you own a Wearable?* *Continued*

	Too complex	Responses in this category indicate that a wearable device was not owned because it was deemed to be too confusing, difficult or complex.	<i>“Don’t understand them”</i>
	Prone to breaking	Responses in this category indicate that a wearable device was not owned because such products break frequently.	<i>“Kept breaking”</i>
	Not allowed to wear	Responses in this category indicate that a wearable device was not owned because the participant is not allowed to wear such products at work.	<i>“Can’t wear watch to work”</i>
	Privacy concerns	Responses in this category indicate that a wearable device was not owned as the participant possessed privacy concerns.	<i>“Don’t want to track myself (privacy concerns)”</i>
	Not applicable to response	Responses in this category are excessively concise or provide little rationale for their response. It might appear as if the reason supports the opposite answer.	<i>“Don’t have it”</i>

Table A.5: Do you own a Smart Appliance?*

Theme	Subtheme	Subsubtheme	Definition	Quote
Yes	Type of device	Smart TV	Responses in this category indicate that a smart appliance was owned because they possessed a smart TV.	<i>“We have a smart TV”</i>
	Applications	TV	Responses in this category indicate that a smart appliance was owned because they appreciated TV on demand and catch-up television.	<i>“To watch telly”</i>
		Entertainment	Responses in this category indicate that a smart appliance was owned because it supports leisure such as watching videos or listening to music. This differs from the previous theme in that it does not refer to television.	<i>“Home entertainment”</i>
		Internet access	Responses in this category indicate that a smart appliance was owned because it is helpful for accessing Internet services.	<i>“Easy access to internet”</i>
		Data synching	Responses in this category indicate that a smart appliance was owned because it allows data to be synched across a range of different devices.	<i>“My TV links to all my other gadgets eg phone ect”</i>
		Food storage	Responses in this category indicate that a smart appliance was owned because it allows food to be chilled and stored.	<i>“Chill my perishables”</i>
	Advantages	Convenient/efficient	Responses in this category indicate that a smart appliance was used due to its convenience and efficiency.	<i>“Makes complicated tasks simpler”</i>
		Usable	Responses in this category indicate that a smart appliance was owned because it is usable.	<i>“Ease of use (when it does what I want it to do)”</i>

Do you own a Smart Appliance?* *Continued*

		Save space	Responses in this category indicate that a smart appliance was owned because it saves space in a property.	<i>"More things to do takes up less space"</i>
		Good value	Responses in this category indicate that a smart appliance was owned because it was deemed to be a wise investment.	<i>"It's expensive but I think is a good investment"</i>
Despite concerns		Privacy concerns	Responses in this category indicate that a smart appliance was owned but the participant still possessed privacy concerns.	<i>"It is like a big brother always spying on you"</i>
		Security concerns	Responses in this category indicate that a smart appliance was owned but the participant still possessed security concerns.	<i>"Concerns regarding security of 'smart' devices though"</i>
		Work	Responses in this category indicate that a smart appliance was owned to undertake work related to their job.	<i>"Work"</i>
		Received	Responses in this category indicate that a smart appliance was owned because it was either received from another person or possessed by a family member.	<i>"Bought as a gift"</i>
		Included as default	Responses in this category indicate that a smart appliance was owned because the smart functionality was included in their product by default.	<i>"Came with this technology built in"</i>
		Not applicable to response	Responses in this category are excessively concise or provide little rationale for their response. It might appear as if the reason supports the opposite answer.	<i>"Na"</i>
Unsure		No need	Responses in this category indicate that they are unsure whether they own a smart appliance, since they have little need for such a system.	<i>"I don't need any of these appliances"</i>
		Not applicable to response	Responses in this category are excessively concise or provide little rationale for their response. It might appear as if the reason supports the opposite answer.	<i>"Why?"</i>
No	Concerns	Privacy	Responses in this category indicate that a smart appliance was not owned as the participant possessed privacy concerns.	<i>"Huge privacy issues, especially with smart TVs"</i>
		Security	Responses in this category indicate that a smart appliance was not owned as the participant possessed security concerns.	<i>"Smart TV is ok but most are major security risks"</i>
	Unrequired	No need	Responses in this category indicate that a smart appliance was not owned because the participant did not need such a device.	<i>"Because I don't need it"</i>
		Objects to technology	Responses in this category indicate that a smart appliance was not owned because the participant dislikes this type of technology.	<i>"It's a gimmick"</i>

Do you own a Smart Appliance?* *Continued*

Disadvantages	Too expensive	Responses in this category indicate that a smart appliance was not owned because it was deemed too expensive.	<i>"Too expensive for my standard"</i>
	Underdeveloped	Responses in this category indicate that a smart appliance was not owned because it was deemed to not be functional enough.	<i>"It's not ready for prime time"</i>
	Too complex	Responses in this category indicate that a smart appliance was not owned because it was deemed to be too confusing, difficult or complex.	<i>"Don't like the added complexity"</i>
	Poor usability	Responses in this category indicate that a smart appliance was not owned because it was deemed to have poor usability.	<i>"The UI is usually terrible on smart TVs"</i>
	Obsolescence	Responses in this category indicate that a smart appliance was not owned because the participant preferred replaceable components. In this manner, add-ons can be exchanged and updated without discarding the appliance when it loses support.	<i>"I would rather have a 'dumb' fridge or TV that can be complemented by more replaceable 'smart' devices"</i>
	Too large	Responses in this category indicate that a smart appliance was not owned because these products are either considered too large or the participant does not have enough space.	<i>"No space"</i>
Prefer laptop		Responses in this category indicate that a smart appliance was not owned because the participant prefers using a laptop.	<i>"Why do I need a smart TV when I have a laptop?!"</i>
Unfamiliar		Responses in this category indicate that a smart appliance was not owned because the technology is still not well-known.	<i>"I've never heard of a smart fridge - I'm uncertain of the benefits"</i>
Upgrade not due		Responses in this category indicate that a smart appliance was not owned because the participant was not ready to upgrade from an older device.	<i>"No need to upgrade my current appliances"</i>
Do not own property		Responses in this category indicate that a smart appliance was not owned because the participant is renting their property. Since they do not own their house, they are limited in the devices they can install.	<i>"Landlord would also be against installing my own fridge"</i>
Not applicable to response		Responses in this category are excessively concise or provide little rationale for their response. It might appear as if the reason supports the opposite answer.	<i>"No particular reason"</i>

Table A.6: Do you own a Home Automation System?*

Theme	Subtheme	Subsubtheme	Definition	Quote
Yes	Type of device	Smart lights	Responses in this category indicate that a smart home device was owned because they possessed smart lights.	<i>"Smart lights at home are good"</i>
	Applications	Communication	Responses in this category indicate that a smart home device was owned because it is helpful for communicating with others.	<i>"Communication and safety"</i>
		Security/safety	Responses in this category indicate that a smart home device was owned because it granted the participant greater security or safety.	<i>"Security"</i>
		Alarm clock	Responses in this category indicate that a smart home device was owned because it is useful as an alarm clock.	<i>"Pretty good for using as an alarm in the morning"</i>
	Advantages	Useful	Responses in this category indicate that a smart home device was owned because it was deemed to provide useful features.	<i>"Convenience"</i>
		Usable	Responses in this category indicate that a smart home device was owned because it is usable.	<i>"Usability"</i>
		Cheap	Responses in this category indicate that a smart home device was owned because the product is not too expensive.	<i>"Acceptable price"</i>
		Fashionable	Responses in this category indicate that a smart home device was owned because it was considered fashionable, desirable or 'cool'.	<i>"Cool factor"</i>
		Saves money	Responses in this category indicate that a smart home device was owned because it saves the participant money.	<i>"I use them because these security systems allow to reduce my mortgage loan"</i>
	Unsure		Responses in this category indicate that a smart home device was owned but the participant is unsure about the reason.	<i>"Unsure"</i>
Not applicable to response		Responses in this category are excessively concise or provide little rationale for their response. It might appear as if the reason supports the opposite answer.	<i>"No"</i>	
Unsure	Not applicable to response		Responses in this category are excessively concise or provide little rationale for their response. It might appear as if the reason supports the opposite answer.	<i>"N/A"</i>
No	Despite concerns	Privacy	Responses in this category indicate that a smart home device was not owned as the participant possessed privacy concerns.	<i>"Does not respect privacy"</i>
		Security	Responses in this category indicate that a smart home device was not owned as the participant possessed security concerns.	<i>"Too easily hacked"</i>

Do you own a Home Automation System?* *Continued*

Unrequired	No need	Responses in this category indicate that a smart home device was not owned because the participant did not need such a device.	<i>"I have no need for it"</i>
	Object to technology	Responses in this category indicate that a smart home device was not owned because the participant dislikes this type of technology.	<i>"Get off your bum and do it yourself"</i>
Disadvantages	Too expensive	Responses in this category indicate that a smart home device was not owned because it was deemed too expensive.	<i>"Probably too expensive"</i>
	Underdeveloped	Responses in this category indicate that a smart home device was not owned because it was deemed to not be functional enough.	<i>"The technology is also not mature"</i>
	Too complex	Responses in this category indicate that a smart home device was not owned because it was deemed to be too confusing, difficult or complex.	<i>"Difficulty of implementation"</i>
Unfamiliar		Responses in this category indicate that a smart home device was not owned because the technology is still not well-known.	<i>"I haven't hear much about it"</i>
Have not taken time yet		Responses in this category indicate that a smart home device was not owned because the participant has not yet taken the time to purchase it.	<i>"Something that I would like to try but haven't so far"</i>
Do not own property		Responses in this category indicate that a smart home device was not owned because the participant is renting their property. Since they do not own their house, they are limited in the systems they can install.	<i>"Would be against my tenancy agreement to install it"</i>
Unsure		Responses in this category indicate that a smart home device was not owned but the participant is unsure of their reason for this decision.	<i>"Don't know"</i>
Not applicable to response		Responses in this category are excessively concise or provide little rationale for their response. It might appear as if the reason supports the opposite answer.	<i>"Never bought one"</i>

A.2 Chapter 6: Device Interviews

Table A.7: Why do you think some people don't buy your device?

Theme	Subtheme	Subsubtheme	Definition	Quote
Form factor	Too large	Too big	Responses in this category indicate that some people might not purchase the device because it is deemed too large or bulky.	<i>"Wouldn't buy one because it's too big"</i>
		Large charger	Responses in this category indicate that some people might not purchase the device since its charger is considered too large.	<i>"It's annoying to having to carry a big bulky charger with you"</i>
	Required peripherals		Responses in this category indicate that some people might not purchase the device as it requires additional peripherals to provide functionality.	<i>"Also because you have to wear a strap as well"</i>
Poor usability	Confusing	Too complex	Responses in this category indicate that some people might not purchase the device as its interface is complex or confusing.	<i>"I suspect that potential end user might be put off by the perceived complexity"</i>
		Few instructions	Responses in this category indicate that some people might not purchase the device as it has/provides few instructions. Therefore, the participant has little support in using the product.	<i>"There's nothing in the manual"</i>
	General usability		Responses in this category indicate that some people might not purchase the device since it has poor usability.	<i>"It's quite small and for someone with clumsy fingers it might be quite hard"</i>
Software issues	Insufficient functionality		Responses in this category indicate that some people might not purchase the device since it does not possess enough useful functionality.	<i>"It may be that they prefer tablets with a wider range of facilities"</i>
	Inaccurate readings		Responses in this category indicate that some people might not purchase the device since its software displays inaccurate data.	<i>"It seems like some Fitbits have a much higher precision when it comes to GPS"</i>
	Compatibility issues		Responses in this category indicate that some people might not purchase the device since it has compatibility issues with other products/manufacturers.	<i>"And then there's compatibility issues around that as well"</i>
Unrequired	Needless	No need	Responses in this category indicate that some people might not purchase the device because they have no need for its features or functionality.	<i>"Because they feel like there's no need"</i>

Why do you think some people don't buy your device? *Continued*

		Not interested in exercise	Responses in this category indicate that some people might not purchase the device because they are not interested in fitness or exercise.	<i>"They are not interested at all about getting a Fitbit or knowing how many steps they're doing"</i>
		Fear of device	Responses in this category indicate that some people might not purchase the device because they are fearful or apprehensive of such products.	<i>"Like my mum, are a little bit afraid of technology"</i>
		Already have similar device	Responses in this category indicate that some people might not purchase the device such they already possess similar products.	<i>"Maybe their TV is fine"</i>
		Unfamiliar	Responses in this category indicate that some people might not purchase the device because they are unfamiliar with such products.	<i>"I think perhaps people don't know enough about them"</i>
Fashion		Dislike the brand	Responses in this category indicate that some people might not purchase the device since they dislike the brand/manufacturer.	<i>"A lot of people hate Apple, firstly"</i>
		Too masculine	Responses in this category indicate that some people might not purchase the device because they regard it as being too masculine. This might deter a woman or somebody who dislikes this style of product.	<i>"So I feel like it's targeting men"</i>
		Prefer non-smart	Responses in this category indicate that some people might not purchase the device since they prefer non-smart functionality.	<i>"Businessmen would rather have a proper fancy Rolex than a sportswatch"</i>
Expense		Expensive	Responses in this category indicate that some people might not purchase the device since they deem it to be too expensive or a waste of money.	<i>"They probably don't want to spend the money on them"</i>
		Poor value	Responses in this category indicate that some people might not purchase the device since they consider it to be poor value.	<i>"Probably get better value for money buying a big box"</i>
Risks	Virtual	Privacy	Responses in this category indicate that some people might not purchase the device since they are concerned over potential privacy risks.	<i>"Because they think that I might be compromised in one way, and might impact their privacy"</i>
		Security	Responses in this category indicate that some people might not purchase the device due to concerns over potential security risks.	<i>"I'm not really sure whether it's a secure piece of equipment or not"</i>

Why do you think some people don't buy your device? *Continued*

	Health	Responses in this category indicate that some people might not purchase the device due to health concerns.	<i>"I've been questioning now like what impact it's been having on my body"</i>
	Familiar with other operating systems	Responses in this category indicate that some people might not purchase the device since they have greater familiarity with other operating systems. Therefore, the challenge from switching might not be warranted.	<i>"Just Windows which is easier for her to get used to"</i>
	Not powerful enough	Responses in this category indicate that some people might not purchase the device since they do not deem it to be fast, quick or powerful enough.	<i>"Because if you want to do stuff that is a bit heavier then maybe a laptop isn't the best machine"</i>
	Constant exercise reminder	Responses in this category indicate that some people might not purchase the device since they do not wish to be frequently reminded to undertake exercise.	<i>"They don't do a lot of exercise so having a constant reminder on their wrist telling them to get up and move"</i>
	Poor connection to phone	Responses in this category indicate that some people might not purchase the device since its connection to a smartphone is poor and unstable.	<i>"Maybe because the quality of connection isn't as good"</i>
	Environmental concerns	Responses in this category indicate that some people might not purchase the device since they believe it might be a waste of materials.	<i>"Waste of materials, I don't know"</i>
Prefer other devices	Prefer smartphones	Responses in this category indicate that some people might not purchase the device since they prefer using smartphones.	<i>"iPhone or whatever other smartphone device"</i>
	Prefer tablets	Responses in this category indicate that some people might not purchase the device since they prefer using tablets.	<i>"And other people that would go for a tablet"</i>
	Prefer desktops	Responses in this category indicate that some people might not purchase the device since they prefer using desktops.	<i>"And I guess otherwise people like desktops"</i>
	Prefer laptops	Responses in this category indicate that some people might not purchase the device since they prefer using laptops.	<i>"I think a lot of people like the traditional laptop style"</i>

Table A.8: How often do you completely turn off your device?

Theme	Subtheme	Subsubtheme	Definition	Quote
Daily/frequently	Energy	Save device battery	Responses in this category indicate that the device is turned off daily/frequently to save the device's battery.	<i>"To save the battery"</i>
		Save energy	Responses in this category indicate that the device is turned off daily/frequently to save energy.	<i>"I don't know, it's just an efficiency thing"</i>
	Ensure performance		Responses in this category indicate that the device is turned off daily/frequently since this contributes to good performance.	<i>"I think it's good for it"</i>
	Habit		Responses in this category indicate that the device is turned off daily/frequently as part of a habit/routine.	<i>"It's just a habit I think"</i>
	Disturbs sleep		Responses in this category indicate that the device is turned off daily/frequently since it disturbs the participant's sleep when the product is left on.	<i>"I keep mine off, it may be better at night when you sleep"</i>
	To receive software updates		Responses in this category indicate that the device is turned off daily/frequently since that allows software updates to take place.	<i>"It means that I can restart it if it has any updates or anything"</i>
	Safety		Responses in this category indicate that the device is turned off daily/frequently since they wish electricity to be disconnected and their products to be safe.	<i>"I don't like the thought of things, I have a dog so I'm worried he'd bite a cable or something"</i>
Between daily and weekly	Convenience		Responses in this category indicate that the device is turned off weekly (or less than daily) since this is more convenient (and requires less effort) than doing it daily/frequently.	<i>"Convenience I think just to pick it up immediately."</i>
	Enable data collection		Responses in this category indicate that the device is turned off weekly (or less than daily) since this allows the product to collect data in the week-long periods in between.	<i>"If I'm using it to track my sleep or to get an understand of my sleep cycle I'm not going to turn it off at night"</i>
	Habit		Responses in this category indicate that the device is turned off weekly (or less than daily) as part of a habit/routine.	<i>"Just habit"</i>
	Storage		Responses in this category indicate that the device is turned off weekly (or less than daily) as the product is sometimes stored to keep it safe.	<i>"It's more safe to switch it off and put it in a drawer"</i>
Between weekly and monthly	Convenience		Responses in this category indicate that the device is turned off monthly (or less than weekly) since this is more convenient (and requires less effort) than doing it daily/frequently or weekly.	<i>"It just doesn't seem necessary to shut it down and wait"</i>

How often do you completely turn off your device? *Continued*

	To provide functionality	Responses in this category indicate that the device is turned off monthly (or less than weekly) since this allows overnight jobs to run in between.	<i>“Usually because I have tasks that are sort of still running overnight”</i>
	Reconfiguration required when restarted	Responses in this category indicate that the device is turned off monthly (or less than weekly) since settings have to be changed each time it is restarted.	<i>“Every time I shut it down and turn it on again I have to reset the BIOS for it to find the hard disk again”</i>
	Good sleep mode	Responses in this category indicate that the device is turned off monthly (or less than weekly) since its sleep mode ensures the product functions correctly.	<i>“Because it’s quite good at sleeping”</i>
	Good battery life	Responses in this category indicate that the device is turned off monthly (or less than weekly) as it has good battery life.	<i>“The battery length is adequate for me”</i>
Never/rarely	Convenience	Responses in this category indicate that the device is turned off rarely or never since this is more convenient (and requires less effort) than doing it more often.	<i>“And yeah, it’s just a bit of a hassle to turn it on and off”</i>
	Enable data collection	Responses in this category indicate that the device is turned off rarely or never since this allows the product to collect data almost constantly.	<i>“I just think if I’m going to count the steps I might as well have it on all the time”</i>
	To provide functionality	Responses in this category indicate that the device is turned off rarely or never as the product must remain on to provide functionality.	<i>“Well because it’s controlling the central heating and hot water in the house, it’s winter, you know, I just leave it running”</i>
	Challenging to turn off	Responses in this category indicate that the device is turned off rarely because it is challenging to undertake this action.	<i>“In order to turn it off, you have to like switch something around at the back”</i>
	Unsure how	Responses in this category indicate that the device is turned off rarely or never since the participant does not know how to undertake this action.	<i>“I don’t actually know how to turn it off”</i>
	Never considered it	Responses in this category indicate that the device is turned off rarely or never since the participant never thought of doing it.	<i>“I’ve never thought about turning it off”</i>

Table A.9: Who do you think has access to your device's data?

Theme	Subtheme	Subsubtheme	Definition	Quote
Hardware vendors	Device manufacturers		Responses in this category indicate that the device's data might be accessible by the main manufacturer of the hardware.	<i>"Probably hardware manufacturers due to the software delivered with it"</i>
	Processor manufacturer		Responses in this category indicate that the device's data might be accessible by the manufacturer of the product's processor.	<i>"Maybe Intel because it's got an Intel processor"</i>
	Audio manufacturer		Responses in this category indicate that the device's data might be accessible by the manufacturer of the product's audio components.	<i>"Realtek audio maybe, they take some stuff?"</i>
Software vendors	Direct	Operating system vendor	Responses in this category indicate that the device's data might be accessible by the main developer (company) of the product's operating system.	<i>"Apple probably get a lot of the diagnostics back"</i>
		App/software vendors	Responses in this category indicate that the device's data might be accessible by the developers of the product's software/apps.	<i>"So if you're using Netflix then probably Netflix that has the data"</i>
		Browser developers	Responses in this category indicate that the device's data might be accessible by the developers of the product's web browser.	<i>"Mozilla probably get a good few things so off Firefox"</i>
		Cloud services	Responses in this category indicate that the device's data might be accessible by a participant's cloud provider or service. This might be through syncing data to online storage.	<i>"I store all of my data on iCloud"</i>
		Google	Responses in this category indicate that the device's data might be accessible by Google.	<i>"Google would get a lot of data from that"</i>
	Shared with third parties		Responses in this category indicate that the device's data might be accessible by third parties which have received it from other entities.	<i>"I'm sure that somehow they are outsourcing the data to somebody else"</i>
Online platforms	Via accounts		Responses in this category indicate that the device's data might be accessible by online platforms through user accounts.	<i>"If I go on Facebook, Facebook has access to it"</i>
	Via cookies		Responses in this category indicate that the device's data might be accessible by online platforms through sessions and cookies.	<i>"There's cookies, most of the websites you visit"</i>
Governments			Responses in this category indicate that the device's data might be accessible by governments (and state actors) through surveillance techniques.	<i>"I would say the government could access things"</i>
Networks	LAN	Company network	Responses in this category indicate that the device's data might be accessible to a participant's company if they are using its network.	<i>"I connect to Private Virtual Network and my employer can look at my data"</i>

Who do you think has access to your device's data? *Continued*

	Networks connected to	Responses in this category indicate that the device's data might be accessible to non-work networks to which the product connects.	<i>"I have used it in airports and public places"</i>
	Internet Service Provider	Responses in this category indicate that the device's data might be accessible by the participant's Internet service provider.	<i>"Possibly my Internet provider"</i>
Other owned devices	Smartphone	Responses in this category indicate that the device's data might be accessible to the participant's smartphone.	<i>"Because it's on my phone"</i>
	Laptop	Responses in this category indicate that the device's data might be accessible the participant's laptop.	<i>"It's only synched up to my laptop"</i>
People	The participant	Responses in this category indicate that the device's data might be accessible to the participant themselves.	<i>"Some of the data it might be just me"</i>
	Family and friends	Responses in this category indicate that the device's data might be accessible to the participant's family, friends and partner.	<i>"My husband only"</i>
	People they choose to share it with	Responses in this category indicate that the device's data might be accessible to those people the participant chooses to share it with.	<i>"Apart from what I give voluntary basically to an online realm"</i>
Phone network provider		Responses in this category indicate that the device's data might be accessible to their smartphone's network provider.	<i>"Possibly your phone provider?"</i>
Hackers		Responses in this category indicate that the device's data might be accessible to hackers or cyber criminals.	<i>"Hence anybody who hacked my WiFi router"</i>
Banks		Responses in this category indicate that the device's data might be accessible to the participant's bank.	<i>"I have my credit card on it so the bank"</i>
Unsure		Responses in this category are unsure who can access the device's data.	<i>"I've no idea"</i>
Never thought about it		Responses in this category indicate that the participant has never considered the matter before.	<i>"I've not really thought about who accesses my data"</i>

Table A.10: Concern Questions - Data Deletion Scenario

Theme	Subtheme	Subsubtheme	Definition	Quote
Concerns	Data	Valued	Responses in this category indicate that the participant opposes their device's data being deleted because it is valued or considered important.	<i>"I have a lot of stuff on there that I wouldn't want to lose"</i>
		Reliant	Responses in this category indicate that the participant opposes their device's data being deleted as they are dependent on it.	<i>"Because I rely on it being there"</i>
	Recovery	Effort to recover	Responses in this category indicate that the participant opposes their device's data being deleted because it would take great effort to recover it.	<i>"I would be mad as hell as I'd have to put it back in"</i>
		Cannot recover data	Responses in this category indicate that the participant opposes their device's data being deleted because it would be impossible to recover it all.	<i>"Some of it I mean I wouldn't be able to get back again"</i>
	Principle	Principles	Responses in this category indicate that the participant opposes their data being deleted on principle.	<i>"I'd be annoyed at the principle of someone doing it"</i>
		Unconsented	Responses in this category indicate that the participant opposes their device's data being deleted without their consent or permission.	<i>"I'd be more frustrated that they did it without my consent"</i>
		Invested in device	Responses in this category indicate that the participant opposes their device's data being deleted since they spent money purchasing the product.	<i>"Just, well I guess I've paid a lot of money for it"</i>
		Feels like theft	Responses in this category indicate that the participant opposes their device's data being deleted as it feels like their information has been stolen.	<i>"It's like someone stealing something from you"</i>
		Security concern	Responses in this category indicate that the participant opposes data deletion since it would make them concerned about the security of their device.	<i>"Well I'd probably feel less secure"</i>
		Lost functionality	Responses in this category indicate that the participant opposes their device's data being deleted as they would lose access to the device's functionality.	<i>"Because of the sleep tracking data which would have like accumulated"</i>
		Inconvenient	Responses in this category indicate that the participant opposes their device's data being deleted as it is inconvenient or annoying.	<i>"It would be very inconvenient"</i>

Concern Questions - Data Deletion Scenario *Continued*

In mitigation	Contingent	Depends what data	Responses in this category indicate that concerns depended on what data was deleted. If it was less important, it would trigger less concern.	<i>“Depends how much data”</i>
		Depends how/why deleted	Responses in this category indicate that concerns depended on how/why the data was deleted. If it occurred due to an accident, the participant would be less worried.	<i>“It probably depends on why they did that”</i>
	Innocuous data		Responses in this category indicate that concerns were limited since the device's data was considered innocuous or not important.	<i>“It's not particularly personal or valuable information”</i>
	Backed-up		Responses in this category indicate that concerns were limited since the device's data was backed up.	<i>“I have backups, I wouldn't be fussed”</i>
	Simple to replace		Responses in this category indicate that concerns were limited since the device's data would be simple to replace. This refers to data being recreated, not restored from a back-up.	<i>“Well because you could just wear it and start again”</i>
	No long-term requirement		Responses in this category indicate that concerns were limited since the device's data was not required over a long term.	<i>“I'm using it more for day-to-day not long-term tracking”</i>
	Other matters are more important		Responses in this category indicate that concerns were limited since the participant finds other matters more important.	<i>“To lose one's settings on their Apple Watch isn't the biggest deal in the world”</i>

Table A.11: Concern Questions - Data Sharing Scenario

Theme	Subtheme	Subsubtheme	Definition	Quote
Concerns	Principle	Principles	Responses in this category indicate that the participant opposes sharing due to their principles or right to privacy/property.	<i>“I think the principle of it would be sort of annoying”</i>
		Unconsented	Responses in this category indicate that the participant opposes their data being shared without their consent or permission.	<i>“But it would be nice if they asked first”</i>
		Private person	Responses in this category indicate that the participant opposes sharing since they consider themselves a private or sensitive individual.	<i>“I'm quite a private person normally”</i>
	Data usage	Tracking	Responses in this category indicate that the participant opposes sharing since it could enable tracking (possibly physical).	<i>“It's indicative of when I'm at home”</i>

Concern Questions - Data Sharing Scenario *Continued*

		Monetisation	Responses in this category indicate that the participant opposes sharing since companies could make money off the data.	<i>"There's a possibility of your data being monetised in some way"</i>
		Unclear motives	Responses in this category indicate that the participant opposes sharing as the motives of the device's company are not known.	<i>"For a purpose I don't know about"</i>
		Valued data	Responses in this category indicate that the participant opposes sharing since they value their device's data.	<i>"I guess a lot of the work stuff I would care"</i>
		Cannot recover data	Responses in this category indicate that the participant opposes sharing since the data cannot be recovered or reclaimed after it has been released.	<i>"Sharing means you can't really control how far it goes"</i>
		Inconvenience	Responses in this category indicate that the participant opposes sharing since it is inconvenient or annoying.	<i>"That would be very inconvenient"</i>
		Company failure	Responses in this category indicate that the participant opposes sharing since they expect the device's company to protect their information.	<i>"But more like you should have been looking after my recipes"</i>
		Unsure	Responses in this category indicate that the participant was unsure of their concerns.	<i>"Like, all these questions I've not actually thought about"</i>
In mitigation	Contingent	Depends what data	Responses in this category indicate that concerns depended on what data was shared. If it was less important, it would trigger less concern.	<i>"It would depend what data it was I think"</i>
		Depends why/how shared	Responses in this category indicate that concerns depended on why/how the data was shared. If it occurred due to an accident, the participant would be less worried.	<i>"It obviously depends what they do with the data"</i>
	Benefits	User benefits	Responses in this category indicate that concerns were limited since the participant receives benefits from their data being shared.	<i>"But if they could look at it and it could help somehow, it doesn't bother me that much"</i>
		Company benefits	Responses in this category indicate that concerns were limited since the device's vendor receives benefits from sharing data.	<i>"Something inane like that which helps Intel make better stuff then okay I can deal with that"</i>
	Little risk	Innocuous data	Responses in this category indicate that concerns were limited since the device's data was considered innocuous or not important.	<i>"Most of it wouldn't be relevant to anybody as such"</i>
		No received damage	Responses in this category indicate that concerns were limited since the participant has never received damage before.	<i>"But for the most part, I don't notice anything happening"</i>

Concern Questions - Data Sharing Scenario *Continued*

	Social norms	Accepted reality	Responses in this category indicate that concerns were limited since the participant accepts data sharing as a modern technological reality.	<i>"I understand, like the concept of, you know, businesses sharing user information"</i>
		Nothing to hide	Responses in this category indicate that concerns were limited since the participant believes they have no data worth hiding.	<i>"I'm not trying to hide anything"</i>
		Exhibitionism	Responses in this category indicate that concerns were limited since the participant would enjoy having their data shown to others.	<i>"I actually would probably pride myself a bit in what I do"</i>

Table A.12: Concern Questions - Data Monitoring Scenario

Theme	Subtheme	Subsubtheme	Definition	Quote
Concerns	Principle	Principles	Responses in this category indicate that the participant opposes monitoring due to their principles or right to privacy/property.	<i>"Well it's an invasion of one's privacy"</i>
		Unconsented	Responses in this category indicate that the participant opposes their data being monitored without their consent or permission.	<i>"That's a bit dodgy if you haven't authorised them to do it"</i>
		Private person	Responses in this category indicate that the participant opposes monitoring since they consider themselves a private or sensitive individual.	<i>"I have got a little plaster over my, over the video camera"</i>
	Data usage	Tracking/security	Responses in this category indicate that the participant opposes monitoring since it could enable tracking (possibly physical).	<i>"You know when I'm in the house if I'm using my scale, for instance"</i>
		Advertising	Responses in this category indicate that the participant opposes monitoring since data could be used for advertisements or marketing purposes.	<i>"Monitoring my data just for the sake of it to advertise things"</i>
		Valued/sensitive data	Responses in this category indicate that the participant opposes monitoring since they value their device's data. This data might also be sensitive or important.	<i>"They can probably get quite a lot of information from someone wearing a Fitbit"</i>
		Uncomfortable / chilling effects	Responses in this category indicate that the participant opposes monitoring as it would make them feel uncomfortable.	<i>"I think it would make me feel very uncomfortable"</i>
		Take action in future	Responses in this category indicate that the participant would take action to protect their data if they felt it was being monitored.	<i>"Next time I would probably be more cautious"</i>

Concern Questions - Data Monitoring Scenario *Continued*

In mitigation	Contingent	Depends what data	Responses in this category indicate that concerns depended on what data was monitored. If it was less important, it would trigger less concern.	<i>“Then that’s more personal so I wouldn’t like that being shared”</i>
		Depends why/how monitored	Responses in this category indicate that concerns depended on why/how the data was monitored. If it occurred due to an accident, the participant would be less worried.	<i>“Yeah again it would depend on who it was and why they were doing it”</i>
	Social norm	Accepted reality	Responses in this category indicate that concerns were limited since the participant accepts data monitoring as a modern technological reality.	<i>“I realise that in the modern age we live in”</i>
		Nothing to hide	Responses in this category indicate that concerns were limited since the participant believes they have no data worth hiding.	<i>“I have nothing to hide so I wouldn’t be bothered”</i>
		Law enforcement	Responses in this category indicate that concerns were limited since the participant believes data should be legitimately monitored for law enforcement purposes.	<i>“If this is something which is done to keep the general population safe...it wouldn’t bother me so much”</i>
	Innocuous data	Responses in this category indicate that concerns were limited since the device’s data was considered innocuous or not important.	<i>“The data itself is not particularly damaging”</i>	
	Functionality benefits	Responses in this category indicate that concerns were limited since the participant receives benefits from their data being monitored.	<i>“It’s the kind of Faustian pact you have every time you Google something, you know”</i>	
	Machine automated	Responses in this category indicate that concerns were limited since their data is being monitored by a machine, rather than a real person.	<i>“If it was a robot that was monitoring what I was doing then I expect I’d be okay with it”</i>	

Table A.13: Concern Questions - Data Selling Scenario

Theme	Subtheme	Subsubtheme	Definition	Quote
Concerns	Principle	Principles	Responses in this category indicate that the participant opposes selling due to their principles or right to privacy.	<i>“I think that would be quite intrusive and violating”</i>
		Unconsented	Responses in this category indicate that the participant opposes data being sold without their consent or permission.	<i>“I didn’t give them permission to do”</i>

Concern Questions - Data Selling Scenario *Continued*

		Intention	Responses in this category indicate that the participant opposes data selling since it is clearly intentional. This contrasts with data sharing, which might be accidental.	<i>"Yeah, so there's the intention too I think"</i>
Property		Ownership	Responses in this category indicate that the participant opposes data selling since they consider themselves to own the information.	<i>"Because it's my stuff, so. It's not yours to sell"</i>
		Bought device	Responses in this category indicate that the participant opposes data selling since they already paid for the device. Therefore, there is no reason for the vendor to need to sell data.	<i>"Because I've bought the Fitbit, I don't think you should sell my data"</i>
Proceeds		Profiting	Responses in this category indicate that the participant opposes data selling since they do not want others profiting from them.	<i>"Because they would be making money"</i>
		No proceeds	Responses in this category indicate that the participant opposes data selling as they expect proceeds from the transactions.	<i>"Pretty upset I think as I want some of that money as well"</i>
Data usage		Security risk	Responses in this category indicate that the participant opposes data selling since it could place them at risk.	<i>"Whether it will allow someone to identify me"</i>
		Advertising	Responses in this category indicate that the participant opposes data selling since information could be used for advertisement or marketing purposes.	<i>"The idea of Facebook monetising something and sending me advertisements"</i>
		Valued/sensitive data	Responses in this category indicate that the participant opposes data selling since they value their device's data. This data might also be sensitive or important.	<i>"I would be horrified, there's so much personal information on there"</i>
		Uncomfortable	Responses in this category indicate that the participant opposes data selling as it would make them feel uncomfortable.	<i>"I would still feel uncomfortable"</i>
		Unknown recipients	Responses in this category indicate that the participant opposes data selling as they are unaware of who will receive their information.	<i>"But it's that lack of knowledge, not knowing where it's going"</i>
In mitigation	Contingent	Depends what data	Responses in this category indicate that concerns depended on what data was sold. If it was less important, it would trigger less concern.	<i>"But I guess it depends what data it is"</i>
		Depends why/how sold	Responses in this category indicate that concerns depended on why/how the data was sold. If it occurred due to an accident, the participant would be less worried.	<i>"If they could show "oh well...it was in our Terms and Conditions", I suppose there isn't much I could do about it"</i>
		Depends if anonymised	Responses in this category indicate that concerns depended on whether sold data was anonymised. If so, participants are likely to be less worried.	<i>"I'd want it to be anonymised"</i>

Concern Questions - Data Selling Scenario *Continued*

	Depends if vendor	Responses in this category indicate that concerns depended on whether the device vendor sold data. If so, participants are likely to be less concerned.	<i>"I think if the company themselves are doing things then that's "meh"</i>
	Accepted reality	Responses in this category indicate that concerns were limited since the participant accepts data selling as a modern technological reality.	<i>"I guess that does happen and in some cases we do accept it"</i>
	Innocuous data	Responses in this category indicate that concerns were limited since the device's data was considered innocuous or not important.	<i>"As I said, I don't really mind who shares it"</i>
	Would sell own data	Responses in this category indicate that concerns were limited as the participant would sell their own data.	<i>"It belongs to me, I would sell it"</i>
	Additional purpose	Responses in this category indicate that concerns were limited since data is not playing other purposes, and so might as well be sold.	<i>"I'm not doing anything with that data"</i>

Table A.14: Behaviour Questions - Have you set a password?

Theme	Subtheme	Subsubtheme	Definition	Quote
Yes	Security	General protection	Responses in this category indicate that the participant set a password to increase their security. This is in general, rather than referring to their device being physically accessed.	<i>"I think a password is crucial for security reasons"</i>
		Physical access	Responses in this category indicate that the participant set a password to limit physical access to her device. This differs from access after loss/theft, and might concern the actions of colleagues or friends.	<i>"If I do leave it out I don't want my office colleagues to be playing pranks on me"</i>
		Loss/theft	Responses in this category indicate that the participant set a password in case the device is lost or stolen.	<i>"But also if it got lost, it wouldn't be so easily accessible by someone else"</i>
		Easy/quick	Responses in this category indicate that the participant set a password because it is quick and easy to use.	<i>"It's easy to set up a password"</i>
		Parental controls	Responses in this category indicate that the participant set a password because it allows them to limit access by their children.	<i>"I have a daughter and I need to restrict the time that she can use it"</i>
		Save battery	Responses in this category indicate that the participant set a password since the screen lock can save device battery.	<i>"If I'm leaving the room for a while and I want to save the battery"</i>
		Habit	Responses in this category indicate that the participant set a password since it is part of their routine.	<i>"Probably habit is the sort of mechanical steps answer"</i>

Behaviour Questions - Have you set a password? *Continued*

	Requirement	Responses in this category indicate that the participant set a password since it was required by software they wished to install.	<i>"I do, yes, because of the bank"</i>
	Duty to protect work computer	Responses in this category indicate that the participant set a password because they believe they have a duty to protect their work computer.	<i>"I'd feel very much that it's my responsibility to look after any sensitive data"</i>
	Separate space	Responses in this category indicate that the participant set a password since the user accounts allow them to have a private space.	<i>"And it's just the idea of having a carved-out personal space which is yours"</i>
	Sensible	Responses in this category indicate that the participant set a password since they believed it a sensible action to take.	<i>"Common sense, you never know what may happen"</i>
	Assisted by other	Responses in this category indicate that the device had a password since the user was assisted by another person.	<i>"Because my boss set it up for me, he automatically sets it up like that"</i>
	Prompted by device	Responses in this category indicate that the participant set a password since the option was prompted to them by the device interface.	<i>"Prompted you as soon as you opened it to set a password"</i>
	Password by default	Responses in this category indicate that the device came with a default password. This was not set by the participant, and they are not sure whether it can be changed/improved.	<i>"It comes with a password"</i>
No	Not supported	Responses in this category indicate that the participant did not set a password as the feature was not supported by the device.	<i>"No, there's nothing"</i>
	Too much effort	Responses in this category indicate that the participant did not set a password as entering the code was considered too much effort.	<i>"It would be a little bit of a hassle entering a password every time"</i>
	Little perceived risk	Responses in this category indicate that the participant did not set a password as they considered the device to be under little threat.	<i>"The very small chance that someone might steal the data"</i>
	Data considered innocuous	Responses in this category indicate that the participant did not set a password since they considered their data to be innocuous.	<i>"There's not much private data on the TV"</i>
	Have not investigated	Responses in this category indicate that the participant did not set a password since they had not investigated this feature. They might also be unsure whether a password exists.	<i>"It's probably some functionality I have not seen"</i>
	Only for online account	Responses in this category indicate that the participant did not set a password on the device, but one is used on the online account.	<i>"The account has a password"</i>

Table A.15: Behaviour Questions - How much time have you spent reading privacy policies?

Theme	Subtheme	Subsubtheme	Definition	Quote
Does read	Gain awareness	Improve knowledge	Responses in this category indicate that the participant read the device's privacy policy to improve their knowledge of the situation.	<i>"I've done something to gain some awareness"</i>
		Check for bad clauses	Responses in this category indicate that the participant read the device's privacy policy to check for egregious or objectionable clauses.	<i>"See if there's anything ridiculous stand out clauses"</i>
	Expertise		Responses in this category indicate that the participant read the device's privacy policy because they have knowledge of technology.	<i>"Just basically because of my professional background I think"</i>
	Principle		Responses in this category indicate that the participant read the device's privacy policy out of principle. They may feel that they should make the effort to remain informed.	<i>"Oh I make a point to try and read it"</i>
	Concerns		Responses in this category indicate that the participant read the device's privacy policy since they are concerned over how their data might be used.	<i>"My desire to protect my privacy survives in the face of reading that privacy policy"</i>
	Little to read		Responses in this category indicate that the participant read the device's privacy policy because it was not much to read.	<i>"I don't remember it being extensive and I went through all of it"</i>
Little/no reading	Document	Too long	Responses in this category indicate that the participant did not read (or barely read) the device's policy because it was too long.	<i>"Because they're big documents"</i>
		Complex/boring	Responses in this category indicate that the participant did not read (or barely read) the device's policy because it was complex, boring or written in legal terminology.	<i>"Because it's written in legalese and it's very boring"</i>
		Foreign language	Responses in this category indicate that the participant did not read (or barely read) the device's policy because it was written in a non-English language.	<i>"Also it was all in Chinese"</i>
		Similar to others	Responses in this category indicate that the participant did not read (or barely read) the device's policy since they believed it would just be the same as other documents.	<i>"I tend to kind of feel if you've read one then you've read many"</i>
	Makes little difference	Need to use device	Responses in this category indicate that the participant did not read (or barely read) the device's policy since they needed to use the product regardless.	<i>"Because you either accept it or you close the laptop down"</i>

Behaviour Questions - How much time have you spent reading privacy policies? *Continued*

	Little real power	Responses in this category indicate that the participant did not read (or barely read) the device's policy since they have little power or recourse. If they object to a condition, they have little influence over the product's vendor.	<i>"I actually don't think it's going to make much of a difference"</i>
	Know enough already	Responses in this category indicate that the participant did not read (or barely read) the device's policy because they believed they knew enough about privacy without it.	<i>"I don't know really, it's probably, I know enough to set the password"</i>
Trust other parties	Trusted data was safe	Responses in this category indicate that the participant did not read (or barely read) the device's policy because they believed their data was safe and secure.	<i>"I guess I probably trusted the company"</i>
	Trust others would check	Responses in this category indicate that the participant did not read (or barely read) the device's policy because they trusted that the document had been analysed by others.	<i>"So I trust those lawyers to have done their job carefully"</i>
Time consuming / too busy		Responses in this category indicate that the participant did not read (or barely read) the device's policy because it would take too much time.	<i>"I have the perception that it's a very lengthy process"</i>
Little perceived risk		Responses in this category indicate that the participant did not read (or barely read) the device's policy because they have not received damage before.	<i>"I haven't heard of any privacy things being breached"</i>
Data considered innocuous		Responses in this category indicate that the participant did not read (or barely read) the device's policy since they considered their data to be innocuous.	<i>"There's nothing that they could access that would worry me"</i>
More concerned about functionality		Responses in this category indicate that the participant did not read (or barely read) the device's policy as they were more concerned about using their product.	<i>"I was more excited to get it on my wrist"</i>
Social norms		Responses in this category indicate that the participant did not read (or barely read) the device's policy because they believed it was common for people to ignore these documents.	<i>"I just hear a lot of people give it 'yes ok, I agree'"</i>
Did not consider		Responses in this category indicate that the participant did not read (or barely read) the device's policy because they did not think about it.	<i>"I just didn't really think about it"</i>

Behaviour Questions - How much time have you spent reading privacy policies? *Continued*

	Did not know policy location	Responses in this category indicate that the participant did not read (or barely read) the device's policy because they did not know where it was placed.	<i>"There may have been in the paper booklet that came with it"</i>
--	------------------------------	---	---

Table A.16: Behaviour Questions - How much time have you spent configuring privacy settings?

Theme	Subtheme	Subsubtheme	Definition	Quote
Did configure	Advantages	Security	Responses in this category indicate that the participant configured their privacy settings to strengthen their security.	<i>"I tend to have like three-quarters of the way towards secure setting"</i>
		Privacy	Responses in this category indicate that the participant configured their privacy settings to improve their privacy.	<i>"I quite like knowing what my laptop's sending"</i>
	Settings	Easy/quick	Responses in this category indicate that the participant configured their privacy settings since they were quick, simple or easy to use.	<i>"It was quite straightforward"</i>
		Prompted	Responses in this category indicate that the participant configured their privacy settings since they were prompted by the device.	<i>"I think probably because when the Mac prompts you"</i>
		Principle	Responses in this category indicate that the participant configured their privacy settings based on principle. They might believe that it is their responsibility to protect their device.	<i>"It's a matter of principle"</i>
		Tweaking/tinkering	Responses in this category indicate that the participant configured their privacy settings since they enjoy customising their device's options.	<i>"When I first set up anything I tend to like to go through and adjust the dials"</i>
		Habit	Responses in this category indicate that the participant configured their privacy settings based on routine.	<i>"I have a routine in doing that when I have a new machine"</i>
		Synched settings from phone	Responses in this category indicate that the participant configured their privacy settings as they were synched across from their smartphone.	<i>"All those settings have been transferred onto the watch"</i>
		Supported by another person	Responses in this category indicate that the participant configured their privacy settings because another individual supported them through the process.	<i>"I probably talked about it for about half an hour while he was...setting it all up for me"</i>
		No rationale	Responses in this category indicate that the participant configured their privacy settings but provided little reason for this decision.	<i>"Five minutes, five minutes maybe"</i>
	Not configure	Little risk	Little perceived risk	Responses in this category indicate that the participant did not configure their privacy settings since they did not believe the device to be at risk.

Behaviour Questions - How much time have you spent configuring privacy settings? *Continued*

	Trust company	Responses in this category indicate that the participant did not configure their privacy settings as they trusted that they were safe with the vendor.	<i>"Like the Fitbit the company would have put in some very basic security"</i>
	Settings not provided	Responses in this category indicate that the participant did not configure their privacy settings since these options were not supported on the device.	<i>"Well the device itself doesn't really have any privacy settings"</i>
	Too much effort	Responses in this category indicate that the participant did not configure their privacy settings as they believed this action would take too much effort.	<i>"I think I'm being lazy. I was being lazy"</i>
	More concerned about functionality	Responses in this category indicate that the participant did not configure their privacy settings as they were more concerned about using the product.	<i>"I think that's just because we were in a hurry to get it set up"</i>
	Data considered innocuous	Responses in this category indicate that the participant did not configure their privacy settings since they considered their data to be innocuous.	<i>"Because I don't mind if somebody had access to all of the data"</i>
	Not prompted	Responses in this category indicate that the participant did not configure their privacy settings as these options were not prompted by the device's interface.	<i>"It's not something that pops out at you"</i>
	Social norms	Responses in this category indicate that the participant did not configure their privacy settings as they believed many people ignored these options.	<i>"People just skirt through it all"</i>

Table A.17: Paradox Questions - Why do people use devices which place privacy at risk?

Theme	Subtheme	Subsubtheme	Definition	Quote
Benefits	Functionality/convenience		Responses in this category indicate that people use risky devices because of the functionality or convenience that they provide. This might be because they consider it to be a trade-off.	<i>"Convenience. Mostly convenience, I think"</i>
		Centralised data	Responses in this category indicate that people use risky devices because they usefully centralise data and features into a single product.	<i>"They use the device for too much"</i>
		Short-term necessity	Responses in this category indicate that people use risky devices because they often have an urgent short-term need. Even if they recognise the threats, they use the technology to complete an immediate task.	<i>"Sometimes you might have to use a laptop like in a cafe or in a hotel"</i>
Lack of awareness		Do not know of risk	Responses in this category indicate that people use risky devices since they do not recognise they are at risk. They are not aware that privacy can be infringed in this manner.	<i>"Because they don't understand privacy, honestly"</i>

Paradox Questions - Why do people use devices which place privacy at risk? *Continued*

	Does not understand protection		Responses in this category indicate that people use risky devices because they are not aware of how to protect themselves.	<i>“And maybe uneducated of what they can do to protect their privacy”</i>
Little perceived risk	Little vulnerability	Not considered vulnerable	Responses in this category indicate that people use risky devices because they do not consider themselves vulnerable. They might know of the threats, but believe themselves safe or not a target.	<i>“I mean, not thinking of themselves as being a legitimate target”</i>
		Trust companies	Responses in this category indicate that people use risky devices since they trust the product’s vendor.	<i>“Or things were put in place to protect, to protect your data”</i>
	Damage not severe		Responses in this category indicate that people use risky devices since they do not think the damage would be severe.	<i>“They have not experienced any privacy breach which is costly”</i>
	Digital disinhibition / intangibility		Responses in this category indicate that people use risky devices because they lose their inhibitions when interacting with technology. Threats and consequences might not appear as tangible when they exist in a virtual environment.	<i>“I just think it’s a very different mind-set in the digital world”</i>
Lack of sensitivity	Data considered innocuous		Responses in this category indicate that people use risky devices since they believe their data is innocuous or unimportant.	<i>“There’s nothing incriminating”</i>
	Nothing to hide		Responses in this category indicate that people use risky devices since they believe they have nothing to hide.	<i>“Oh well, you know I’ve got nothing to hide”</i>
Lack of effort	Apathy		Responses in this category indicate that people use risky devices because they cannot be bothered to protect themselves.	<i>“Don’t really care, can’t be bothered”</i>
	Fatalism		Responses in this category indicate that people use risky devices since they believe their privacy will be infringed anyway. They might believe that technology companies are too powerful to resist.	<i>“People might even be aware but they just feel like they’re powerless to fix it”</i>
Personality	Extroverted		Responses in this category indicate that people use risky devices because they are an extroverted or public person.	<i>“If you don’t mind then you’re just a public person”</i>
	Optimistic		Responses in this category indicate that people use risky devices because they are optimistic that they will not fall victim.	<i>“Then we all just collectively assume that we are safe”</i>

Paradox Questions - Why do people use devices which place privacy at risk? *Continued*

Do not care about privacy	Responses in this category indicate that people use risky devices because they simply do not care about privacy.	<i>“So I think a lot of people just don’t care”</i>
Easier to use risky devices	Responses in this category indicate that people use risky devices because they find them more usable. This might be because they lack privacy restrictions.	<i>“Or maybe their device is easier to use”</i>
Social norms / pressure	Responses in this category indicate that people use risky devices because there is a social norm to interact with technology. There might also be peer pressure from colleagues to use such products.	<i>“And it kind of became the norm, the use of tablets and smartphones”</i>
Expense	Responses in this category indicate that people use risky devices because it is too expensive to use safer alternatives. A person might be ‘paying’ for a service with their data, and it might be more expensive if that data was not shared.	<i>“And if you do take a cheap Android phone it’s probably configured to leak privacy”</i>
Fake software	Responses in this category indicate that people use risky devices because they might be tricked by fake/fraudulent software.	<i>“Secondly, some people can design very genuine software”</i>

Table A.18: Paradox Questions - Why do people use their devices in an ‘unprivate’ way?

Theme	Subtheme	Subsubtheme	Definition	Quote
	Functionality/convenience		Responses in this category indicate that devices are used in an ‘unprivate’ manner because of the functionality or convenience that their features can provide. Individuals might regard privacy and functionality as a trade-off, with the former being more important.	<i>“Because they want the features”</i>
Lack of awareness	Do not know of risk		Responses in this category indicate that devices are used in an ‘unprivate’ manner since they people do not recognise they are at risk. They are not aware that privacy can be infringed in this manner.	<i>“Again, because they’re not aware of the risks”</i>
	Do not understand protection		Responses in this category indicate that devices are used in an ‘unprivate’ manner because people are not aware of how to protect themselves.	<i>“They don’t know in many cases even how to change the settings”</i>

Paradox Questions - Why do people use their devices in an 'unprivate' way? *Continued*

	Generational differences		Responses in this category indicate that devices are used in an 'unprivate' manner because older people may be less aware of how to use technology.	<i>"My parent's generation, who are not as technologically-literate as we are"</i>
Little perceived risk	Trust companies		Responses in this category indicate that devices are used in an 'unprivate' manner since people trust the product's vendor.	<i>"There's a certain amount of trust you have with the developer"</i>
	Little severity	Little potential damage	Responses in this category indicate that devices are used in an 'unprivate' manner since people do not think the damage would be severe.	<i>"So maybe nothing happened, nothing serious happened to them yet"</i>
		Consequences not immediate	Responses in this category indicate that devices are used in an 'unprivate' manner since potential negative consequences are not immediate. Therefore, it becomes challenging to logically connect the events.	<i>"The consequences are far off and they don't happen immediately"</i>
Lack of sensitivity	Data considered innocuous		Responses in this category indicate that devices are used in an 'unprivate' manner since people believe their data is innocuous or unimportant.	<i>"They don't see the value that someone could gain or obtain from knowing the data"</i>
	Nothing to hide		Responses in this category indicate that devices are used in an 'unprivate' manner since people believe they have nothing to hide.	<i>"Because they have nothing to hide, maybe"</i>
Personality	Extroverted		Responses in this category indicate that devices are used in an 'unprivate' manner because people are extroverted or open.	<i>"They want to boast about the things they've done"</i>
	Denial		Responses in this category indicate that devices are used in an 'unprivate' manner because people deny the fact that they might be under threat.	<i>"Maybe denial to an extent as well"</i>
Device interface	Design		Responses in this category indicate that devices are used in an 'unprivate' manner because of the way the product is designed. Perhaps a user is encouraged to use the device in a way which could threaten their privacy.	<i>"They try quite hard because their business model is to get your information"</i>
	Default settings		Responses in this category indicate that devices are used in an 'unprivate' manner because default settings are often quite lax.	<i>"A lot of the devices default to not being incredibly private"</i>

Paradox Questions - Why do people use their devices in an ‘unprivate’ way? *Continued*

Lack of effort	Apathy	Responses in this category indicate that devices are used in an ‘unprivate’ manner because people cannot be bothered to protect themselves.	<i>“Yeah, I think general apathy”</i>
	Easier to use	Responses in this category indicate that devices are used in an ‘unprivate’ manner because it is easier to use them that way.	<i>“Because it’s easier. Easier, quicker, probably”</i>
Do not care about privacy		Responses in this category indicate that devices are used in an ‘unprivate’ manner because people simply do not care about privacy.	<i>“Or they simply don’t care”</i>
Social norms / pressure		Responses in this category indicate that devices are used in an ‘unprivate’ manner because there is a social norm to interact with technology. There might also be peer pressure from colleagues to use such products.	<i>“And they’re trying to get social gratification from people”</i>
Too expensive		Responses in this category indicate that devices are used in an ‘unprivate’ manner because it is too expensive to use it in a safer way. The person might be ‘paying’ for that service with their data, and it might be more expensive if that data was not shared.	<i>“Would rather use the free voice over internet app rather than the paid version”</i>
Shared device		Responses in this category indicate that devices are used in an ‘unprivate’ manner because a product might be shared by multiple people.	<i>“They might be like a public device, like a family device”</i>

Table A.19: Paradox Questions - Why do people claim to value privacy but use risky devices?

Theme	Subtheme	Subsubtheme	Definition	Quote
Benefits	Functionality/convenience		Responses in this category indicate that people use devices despite their claims because of the functionality or convenience that these products provide. Individuals might regard privacy and functionality as a trade-off, with the former being more important.	<i>“I think convenience”</i>
	Short-term necessity		Responses in this category indicate that people use devices despite their claims because they often have an urgent short-term need. Even if they recognise the threats, they use the technology to complete an immediate task.	<i>“But then on the other hand they have to use computers ... and time pressure”</i>

Paradox Questions - Why do people claim to value privacy but use risky devices? *Continued*

Lack of awareness	Do not know of risk		Responses in this category indicate that people use devices despite their claims since they do not recognise they are at risk. They are not aware that privacy can be infringed in this manner.	<i>“Not having a full awareness of what they’re doing is risky behaviour”</i>
	Do not understand protection		Responses in this category indicate that people use devices despite their claims because they are not aware of how to protect themselves.	<i>“But I have some idea what to do, I don’t think everyone does”</i>
Little perceived risk	Little vulnerability	Not considered vulnerable	Responses in this category indicate that people use devices despite their claims because they do not consider themselves vulnerable. They might know of the threats, but believe themselves safe or not a target.	<i>“Nobody is interested in what I’m doing or whatever”</i>
		Trust companies	Responses in this category indicate that people use devices despite their claims since they trust the product’s vendor.	<i>“You just assume that it’s new and it’s all protected”</i>
		Data considered innocuous	Responses in this category indicate that people use devices despite their claims since they believe their data is innocuous or unimportant.	<i>“The data that I’m putting at risk is so insignificant to me”</i>
	Little severity	Little potential damage	Responses in this category indicate that people use devices despite their claims since they do not think the damage would be severe.	<i>“If the level of risk is not an extreme level”</i>
		Consequences not immediate	Responses in this category indicate that people use devices despite their claims since potential negative consequences are not immediate. Therefore, it becomes challenging to logically connect the events.	<i>“You don’t really think about the consequences because they’re not immediate”</i>
	Digital disinhibition / intangibility		Responses in this category indicate that people use devices despite their claims because they lose their inhibitions when interacting with technology. Threats and consequences might not appear as tangible when they exist in a virtual environment.	<i>“Doing it on your computer now, you just don’t realise just how far it goes”</i>
Lack of effort	Apathy		Responses in this category indicate that people use devices despite their claims because they cannot be bothered to protect themselves.	<i>“I just get lazy about it sometimes, but that’s what happens”</i>
	Fatalism		Responses in this category indicate that people use devices despite their claims since they believe their privacy will be infringed anyway. They might believe that technology companies are too powerful to resist.	<i>“I think I just gave up”</i>

Paradox Questions - Why do people claim to value privacy but use risky devices? *Continued*

Personality	Optimistic		Responses in this category indicate that people use devices despite their claims because they are optimistic that they will not fall victim.	<i>"I tell myself that everything is really safe"</i>
	Group actions	Privacy is subjective	Responses in this category indicate that people use devices despite their claims because privacy itself is a contextual and subjective topic. Privacy might be conceptualised differently by different people or in different environments.	<i>"I think privacy is very ambiguous"</i>
		Human nature	Responses in this category indicate that people use devices despite their claims because it is human nature to desire contrasting goals.	<i>"Good old fashioned hypocrisy"</i>
Society	Social norms / pressure		Responses in this category indicate that people use devices despite their claims because there is a social norm to interact with technology. There might also be peer pressure from colleagues to use such products.	<i>"Perhaps social pressures, it's just like Facebook"</i>
	Generational differences		Responses in this category indicate that people use devices despite their claims because they might have grown up with digital technology. They therefore might accept data sharing in a manner that older generations might have not.	<i>"Feel they are acting in a private way, compared to someone who is in their 40s"</i>
Device design			Responses in this category indicate that people use devices despite their claims because the devices are designed to extract information.	<i>"You set up an Android phone and it will ask you all these details"</i>
Inconvenient			Responses in this category indicate that people use devices despite their claims because it is inconvenient to use products in a privacy-protecting manner.	<i>"You know that you should have a password but ... it's not convenient enough"</i>
Control/agency			Responses in this category indicate that people use devices despite their claims because they control what they actively broadcast. A person might not mind sharing their own data, but would dislike another party accessing it.	<i>"More about thinking about the choice you're making before you do it"</i>
Setting an example to others			Responses in this category indicate that people use devices despite their claims as their claims are used to set a good example.	<i>"If they're a parent then maybe they want to set a good role"</i>

A.3 Chapter 7: Prototype Evaluation

Evaluation questionnaire tables

Table A.20: What did you like most about the game?

Theme	Definition	Quote
Educational/instructive	Responses in this category indicate that the most liked element was the provided education. The user was informed through messages, hints and direct feedback.	<i>"I liked that you were able to learn something"</i>
Enjoyable/amusing	Responses in this category indicate that the most liked element was that the game was fun. The user enjoyed the experience and found the app to be amusing.	<i>"Was kind of fun as well"</i>
Concept	Responses in this category indicate that the most liked element was the concept of the game. The user appreciated the novel smartwatch approach and how it differed from other games.	<i>"That it incorporated the virtual device's settings"</i>
Movement usability	Responses in this category indicate that the most liked element was the ease of movement. The user found the game to be easy to play.	<i>"The ease of moving the character"</i>
Simplicity	Responses in this category indicate that the most liked element was the simplicity of the game. The user found it easy to understand and was not confused by the design.	<i>"Plain and easy, straight to the point"</i>
Challenges	Responses in this category indicate that the most liked element was the privacy challenges (also referred to as quests or tasks).	<i>"I liked trying to do the challenges as fast as possible"</i>
Questions	Responses in this category indicate that the most liked element was the villager questions.	<i>"I liked testing my knowledge when meeting the neighbours"</i>
Collecting coins	Responses in this category indicate that the most liked element was the act of collecting coins. This refers to picking up coins while navigating the map, rather than winning points by completing challenges or questions.	<i>"I liked collecting coins"</i>
Graphics	Responses in this category indicate that the most liked element was the graphics. They might have found the appearance of the game to be cute or friendly.	<i>"The retro graphics"</i>
Reminiscent of old games	Responses in this category indicate that the most liked element was the retro design. The navigation and graphics reminded the user of old games.	<i>"It reminded me of old Nintendo games"</i>
Characters/avatar	Responses in this category indicate that the most liked element was the characters. This might refer to the villagers, thieves or customised avatars.	<i>"Customization of the character"</i>

What did you like most about the game? *Continued*

Map design / navigation	Responses in this category indicate that the most liked element was the map layout. The user enjoyed navigating the town and exploring the landmarks.	<i>"The layout of the town was cool"</i>
Gaining points	Responses in this category indicate that the most liked element was the accruing of points.	<i>"I liked earning points"</i>
Quick	Responses in this category indicate that the most liked element was the brief nature of the game. The app did not take long to play, and hence did not waste time.	<i>"That you could do it fairly quickly"</i>
Interactive	Responses in this category indicate that the most liked element was the interactivity.	<i>"The interactivity"</i>
Unpredictable	Responses in this category indicate that the most liked element was the unpredictability of the game.	<i>"Not know what was going to happen next"</i>
Easy difficulty	Responses in this category indicate that the most liked element was the ease of completion.	<i>"It was very easy"</i>
Challenging difficulty	Responses in this category indicate that the most liked element was the challenging gameplay.	<i>"I found it hard"</i>
Unsure	Responses in this category were unsure what their most liked element was.	<i>"Not sure"</i>
Little/nothing	Responses in this category did not appear to like little of the smartwatch game.	<i>"Nothing at all"</i>
Not applicable to response	Responses in this category are excessively concise or provide little rationale for their response. It might appear as if these replies answer a different question (e.g., dislikes).	<i>"NO"</i>

Table A.21: What did you like least about the game?

Theme	Subtheme	Definition	Quote
Usability issues	General controls	Responses in this category disliked the general usability of the game. They found the clicking to be difficult, particularly in desktop environments.	<i>"The click and drag needed to move the player is very annoying"</i>
	Challenge issues	Responses in this category disliked the usability of the challenges. There might have been certain tasks which were challenging to complete.	<i>"I couldn't navigate back from the privacy tasks to the game"</i>
	Laggy / slow navigation	Responses in this category disliked the slow responsiveness of the game. They may have also opposed the effort required to navigate the game map.	<i>"I wasn't so fond of the lag"</i>

What did you like least about the game? *Continued*

	Unfamiliar interface	Responses in this category disliked the unfamiliarity of the interface. Therefore, they found it difficult to know where to go.	<i>"I couldn't found how to find the settings"</i>
	Bugs/errors	Responses in this category disliked the bugs within the game, which detracted from the general experience.	<i>"Sometimes it did not worked well, i could not move at first click"</i>
Boring	Theme	Responses in this category disliked the general theme of the game, labelling it as 'boring' or 'pointless'.	<i>"It was a bit boring"</i>
	Repetitive	Responses in this category disliked the repetitive nature of the game.	<i>"It was repetitive"</i>
Characters	Thieves	Responses in this category disliked encountering the thief characters.	<i>"I didn't like the thieves"</i>
	Villagers	Responses in this category disliked encountering the villager characters.	<i>"Villagers"</i>
Poor design / graphics		Responses in this category disliked the graphics of the game. This might be because they are deemed too basic.	<i>"Graphics could be better"</i>
Too difficult		Responses in this category disliked the game as they considered it too difficult or challenging to complete.	<i>"How quick your health goes down"</i>
Lack of features		Responses in this category disliked that the game lacked extra features and customisation.	<i>"The lack of features such as awards or attacking"</i>
Too long		Responses in this category disliked the fact that the game was so long in duration.	<i>"Long"</i>
Too short		Responses in this category disliked the fact that the game was so short in duration.	<i>"A bit short, it could have been longer"</i>
Lack of guidance		Responses in this category disliked the fact that the game included little guidance. They would have appreciate hints and assistance for the challenges.	<i>"Lack of any guidance on the 'settings' app"</i>
Preachy		Responses in this category disliked the 'preachy' narrative of the game, which repeatedly emphasised the importance of privacy.	<i>"A bit preachy"</i>
Knew information already		Responses in this category disliked the fact that they knew much of the privacy content already. Therefore, they learned little from the experience.	<i>"I also learned nothing new, since I was already aware of how to protect myself"</i>
No music		Responses in this category disliked the fact that there was no sound in the game.	<i>"No music"</i>

What did you like least about the game? *Continued*

Unclear map	Responses in this category disliked the map's perceived lack of clarity. This might have been because only a portion of the town was visible at any one time.	<i>"I couldn't see whole map where I was going"</i>
Too simplistic	Responses in this category disliked the game since it was considered too simplistic.	<i>"That it seemed super basic"</i>
Coins	Responses in this category disliked the role of coins within the game.	<i>"Coins"</i>
Would require excessive permissions	Responses in this category disliked the fact that a smartwatch game would require a large amount of permissions to support the challenges. It should be noted that since the prototype was hosted online, privacy permissions were not required.	<i>"It could detect out of game actions, seemingly giving it a lot of permissions on the watch"</i>
Made user competitive	Responses in this category disliked the fact that the timed challenges made the user competitive.	<i>"Invigorated the competitive side of me"</i>
Everything	Responses in this category reported disliking everything about the game.	<i>"Everything"</i>
Nothing	Responses in this category disliked little or nothing about the game.	<i>"Nothing I liked it all"</i>
Unsure	Responses in this category were unsure what they disliked most about the game.	<i>"Not sure"</i>
Not applicable to response	Responses in this category are excessively concise or provide little rationale for their response. It might appear as if these replies answer a different question (e.g., likes).	<i>"Not applicable"</i>

Table A.22: If you could improve one thing about the game, what would it be?

Theme	Subtheme	Definition	Quote
Usability of controls		Responses in this category believe the game should be improved by enhancing the usability of the gameplay controls.	<i>"I'd make it easier to move around"</i>
Design/graphics		Responses in this category believe the game should be improved by enhancing the design and graphics.	<i>"Make the graphics a little more this century"</i>
Extend / adjust game	Levels	Responses in this category believe the game should be improved by extending, adjusting or adding more levels.	<i>"More levels"</i>

If you could improve one thing about the game, what would it be? *Continued*

	Challenges	Responses in this category believe the game should be improved by extending, adjusting or adding more challenges/tasks.	<i>“Longer game with more tasks”</i>
	Questions	Responses in this category believe the game should be improved by adjusting the villager questions.	<i>“Make the answers to the privacy questions less obvious”</i>
	Game modes	Responses in this category believe the game should be improved by introducing more modes or adjusting the gameplay.	<i>“More minigames”</i>
	Maps	Responses in this category believe the game should be improved by extending, adjusting or adding more maps.	<i>“You could go inside the buildings for adventures”</i>
Add sound/music		Responses in this category believe the game should be improved by introducing sound effects or music.	<i>“Maybe add some sound to the game”</i>
Speed/responsiveness		Responses in this category believe the game should be improved by enhancing the speed and responsiveness of the interface.	<i>“More fluid movement”</i>
More hints/instructions		Responses in this category believe the game should be improved by providing more help, feedback, hints and instructions.	<i>“Some explanations, when the answer is not correct”</i>
More randomisation		Responses in this category believe the game should be improved by introducing more randomisation into the challenges and the map.	<i>“Make it more random, for example make some of the encounters not be at the same place”</i>
More coins		Responses in this category believe the game should be improved by adding more rewards and coins.	<i>“More coins”</i>
Improve characters		Responses in this category believe the game should be improved by enhancing the characters.	<i>“The good guys”</i>
Apply to different smartwatches		Responses in this category believe the game should be improved by considering smartwatches other than just Wear OS.	<i>“Have an Apple or Android option”</i>
Implement natively on watch		Responses in this category believe the game should be improved by implementing a version on a real smartwatch.	<i>“Actually download an app and play on a watch”</i>
Improved programming		Responses in this category believe the game should be improved by writing the code in a different manner.	<i>“Change the language its written in”</i>
Change/remove challenge timer		Responses in this category believe the game should be improved by adjusting or removing the timer in the challenges.	<i>“Don’t have timed tasks”</i>
Remove introductory video		Responses in this category believe the game should be improved by removing the tutorial video.	<i>“Maybe not show the tutorial fully”</i>

If you could improve one thing about the game, what would it be? *Continued*

Fix current errors	Responses in this category believe the game should be improved by fixing the current bugs and interface errors.	<i>"The button on side didn't work"</i>
Add exciting theme	Responses in this category believe the game should be improved by adjusting the theme.	<i>"More purpose"</i>
Unsure	Responses in this category are unsure how the game should be improved.	<i>"Don't know"</i>
Nothing	Responses in this category do not believe that the game should be adjusted.	<i>"None, the game is perfect"</i>
Everything	Responses in this category believe that the game needs to be totally altered to improve the experience.	<i>"Everything"</i>
Not applicable to response	Responses in this category are too vague, concise or irrelevant to provide advice.	<i>"N/A"</i>

Table A.23: What would influence you to change your behaviour?

Theme	Subtheme	Definition	Quote
Damage	Personally affected/targeted	Responses in this category indicate that behaviour would be changed if the user was personally affected or targeted by an attack.	<i>"If my smart watch information was stolen. Oh gosh"</i>
	Friends affected/targeted	Responses in this category indicate that behaviour would be changed if a user's friend or family was personally affected or targeted by an attack.	<i>"If someone I knew had been affected"</i>
Information	Threats proven/demonstrated	Responses in this category indicate that behaviour would be changed if privacy threats could be clearly proven or demonstrated.	<i>"Seeing real life examples of people being exploited"</i>
	More details	Responses in this category indicate that behaviour would be changed if the user received additional information.	<i>"More information"</i>
	News reports	Responses in this category indicate that behaviour would be changed if the user read news reports or articles about privacy risks.	<i>"More news stories about data theft"</i>
	Friends advice	Responses in this category indicate that behaviour would be changed if a user's friend advised them to take more control over their privacy.	<i>"A friend giving me advice"</i>
	Awareness advertisement	Responses in this category indicate that behaviour would be changed if the user received additional information from awareness campaigns/initiatives.	<i>"Advertising campaigns would also help"</i>

What would influence you to change your behaviour? *Continued*

Features	Current protective features	Responses in this category indicate that behaviour will be changed by the presence of the current privacy settings.	<i>“To be more careful and turn locations off”</i>
	New features	Responses in this category indicate that behaviour would be changed by the addition of new features or settings.	<i>“Additional privacy options that I can use that aren’t currently available”</i>
	Improved usability/convenience	Responses in this category indicate that behaviour would be changed if privacy settings were easier to use or more convenient.	<i>“If the process of changing the privacy options are easy, then it will influence me to change it”</i>
Educational media	Games	Responses in this category indicate that behaviour would be changed if they played additional educational games.	<i>“Playing more games like these”</i>
	Videos	Responses in this category indicate that behaviour would be changed if they watched additional educational videos.	<i>“Educational videos”</i>
Session was persuasive		Responses in this category indicate that behaviour was changed in response to the privacy game or introductory video.	<i>“The game really helped as it was fun as well as making me aware of the risks”</i>
Increased risk		Responses in this category indicate that behaviour would be changed in response to an increased degree of privacy risk.	<i>“The risk will make me change how secure my smartwatch data”</i>
Storing sensitive data		Responses in this category indicate that behaviour would be changed if the user had more sensitive data on their watch.	<i>“A banking app”</i>
Advantages/incentives		Responses in this category indicate that behaviour would be changed if privacy settings offered additional advantages or incentives.	<i>“Rewards for doing it”</i>
Required by apps		Responses in this category indicate that behaviour would be changed if apps required that privacy settings were adjusted.	<i>“If an app required it”</i>
Nothing/little		Responses in this category indicate that nothing or little could be done to further encourage behaviour change.	<i>“Nothing, not that I care much about the privacy”</i>
Unsure		Responses in this category are unsure what could encourage an alteration in behaviour.	<i>“I don’t know right now”</i>
Not applicable to response		Responses in this category either do not relate to the question, are too vague, excessively concise, or do not provide rationale.	<i>“N/A”</i>

Posttest questionnaire tables

Table A.24: Do you think you understand how to protect your smartwatch data?

Theme	Subtheme	Subsubtheme	Definition	Quote
Yes	Learned through session		Responses in this category indicate that they understand protection because they learned it through the session (game or video).	<i>“Done a game a few days ago that thought me how to change settings”</i>
	Learned independently		Responses in this category indicate that they understand protection since they have learned about it through their own efforts.	<i>“I recently read an article on the subject and learned a lot”</i>
	Familiarity/expertise in technology		Responses in this category indicate that they understand protection since they have familiarity or expertise in technology.	<i>“I’m pretty tech savvy”</i>
	Concerned about privacy risk		Responses in this category indicate that they understand protection since they were incentivised by their concern for privacy.	<i>“Because privacy of data is important or else it could be misused by third parties”</i>
	Read information/instructions		Responses in this category indicate that they understand protection since they read information or instructions.	<i>“Because I read the instructions”</i>
	Experience/knowledge of protective settings		Responses in this category indicate that they understand protection since they either have knowledge or experience of the privacy settings.	<i>“I know to disable apps and permissions”</i>
	Similar to other devices		Responses in this category indicate that they understand protection since smartwatches are considered to be similar to other devices. If the user can protect their smartphone, these skills can extend to the watch.	<i>“I feel as though it’s similar to my phone”</i>
	Cautious		Responses in this category indicate that they understand protection since they are cautious with their smartwatch.	<i>“I’m careful with what I download”</i>
	Not applicable to response		Responses in this category are either too concise, too vague, unrelated to the question or not clear in their rationale.	<i>“Because”</i>
Unsure	Insufficiently skilled	Not informed	Responses in this category indicate that they are unsure about protection since they do not know enough about the topic.	<i>“Not well versed in smartwatch security”</i>
		Have not thought to explore	Responses in this category indicate that they are unsure about protection since they have not thought about exploring privacy features.	<i>“I never really thought about it”</i>

Do you think you understand how to protect your smartwatch data? *Continued*

		Smartwatches are novel	Responses in this category indicate that they are unsure about protection since smartwatches are still novel.	<i>“Relatively new technology”</i>
		Too complex	Responses in this category indicate that they are unsure about protection since guarding smartwatch privacy is highly complex.	<i>“They don’t seem as straight forward as phones”</i>
		Adversaries are too skilled	Responses in this category indicate that they are unsure about protection since they believe other parties might be skilled in accessing data.	<i>“There are always new ways to hack being discovered”</i>
		Unconcerned	Responses in this category indicate that they are unsure about protection since they do not care greatly about their privacy.	<i>“I do not care if my data is used by other companies”</i>
		Try but not confident	Responses in this category indicate that they are unsure about protection despite trying to guard their data.	<i>“I think I have a good idea. But there are always new ways”</i>
		Assumed was secure	Responses in this category indicate that they are unsure about protection since they assumed their smartwatch was secure by default.	<i>“I’m not really sure how to protect my smartwatch data. I just assumed it was already secure”</i>
		Protection not required	Responses in this category indicate that they are unsure about protection since they did not believe it was necessary.	<i>“No need protection”</i>
		Do not take action	Responses in this category indicate that they are unsure about protection since they didn’t take steps to guard their data.	<i>“I don’t actively do anything protect it”</i>
		Traded for functionality	Responses in this category indicate that they are unsure about protection since privacy sometimes must be relaxed to allow features.	<i>“The app will not function without these permissions”</i>
		Unsure	Responses in this category cannot explain why they are unsure about smartwatch protection.	<i>“I’m not pretty sure about it”</i>
		Not applicable to response	Responses in this category are either too concise, too vague, unrelated to the question or not clear in their rationale.	<i>“Just because”</i>
No	Insufficiently skilled	Not informed	Responses in this category indicate that they doubt they understand protection since they do not know enough about the topic.	<i>“I really have no idea how to protect my data”</i>
		Have not explored	Responses in this category indicate that they doubt they understand protection since they did not explore the privacy settings.	<i>“Have not taken the time to fully understand all threats”</i>

Do you think you understand how to protect your smartwatch data? *Continued*

	Smartwatches are novel	Responses in this category indicate that they doubt they understand protection since smartwatches are still novel.	<i>"I'm not familiar enough with my smartwatch"</i>
	Too complex	Responses in this category indicate that they doubt they understand protection since guarding privacy is too complex.	<i>"Too many options - seems to vary by app"</i>
	Adversaries are too skilled	Responses in this category indicate that they doubt they understand protection since other parties are skilled at data access.	<i>"Thieves are always coming up with new ways to get into our technology"</i>
	Unconcerned	Responses in this category indicate that they doubt they understand protection since they are not concerned or interested about their privacy.	<i>"Never was a concern"</i>
	Protection not required	Responses in this category indicate that they doubt they understand protection since they did not believe it was required.	<i>"Safe"</i>
	Assumed was secure	Responses in this category indicate that they doubt they understand protection since they assumed that the smartwatch was secure.	<i>"Thought it was safe and protected"</i>
	Lack of instructions	Responses in this category indicate that they doubt they understand protection since sufficient instructions are not available.	<i>"Because there was no instructions"</i>
	Lack of salience	Responses in this category indicate that they doubt they understand protection since the topic was not salient. As they did not think about privacy, it is unlikely they understand it.	<i>"Because I've never really thought about it before"</i>
	Unsure	Responses in this category are unsure why they do not understand smartwatch protection.	<i>"I don't know"</i>
	Not applicable to response	Responses in this category are either too concise, too vague, unrelated to the question or not clear in their rationale.	<i>"N/a"</i>

Table A.25: Do you feel confident you can protect your smartwatch data?*

Theme	Subtheme	Subsubtheme	Definition	Quote
Yes	Learned through session		Responses in this category indicate that they are confident since they learned protection through the session (game or video).	<i>"I have played a game that teaches best practice"</i>
	Learned independently		Responses in this category indicate that they are confident since they learned protection through their own efforts, or through discussion with others.	<i>"I have done a lot of research on the subject"</i>
	Taken action to protect data		Responses in this category indicate that they are confident since have used tools to protect their data.	<i>"I am very careful with my security settings and always do all the updates"</i>
	Understand smartwatch protection		Responses in this category indicate that they are confident since they understand how protection works on a smartwatch.	<i>"Because I am aware of the risks and I know how to minimize it"</i>
	Familiarity/expertise in technology		Responses in this category indicate that they are confident since they have familiarity or expertise in technology.	<i>"I did IT networking and security"</i>
	Similar to other devices		Responses in this category indicate that they are confident since smartwatch protection is deemed similar to protecting other devices.	<i>"I feel it is similar to a phone"</i>
	Cautious		Responses in this category indicate that they are confident since are cautious with their data and are concerned about privacy.	<i>"I am careful with what is installed"</i>
	Not under great threat		Responses in this category indicate that they are confident since they do not believe they are under great threat.	<i>"I think I'm safe"</i>
	Trust vendor		Responses in this category indicate that they are confident since they trust the protection and updates of the watch manufacturer.	<i>"I believe the security futures are good enough"</i>
	Not applicable to response		Responses in this category are either too concise, too vague, unrelated to the question or not clear in their rationale.	<i>"Believe"</i>
Unsure	Insufficiently skilled	Not informed	Responses in this category indicate that they are unsure since they do not know enough about the topic.	<i>"I don't feel like I have enough knowledge on the how to protect my data"</i>
		Too complex	Responses in this category indicate that they are unsure since protecting data is complex.	<i>"Too many variables"</i>
	Infeasible	Adversaries are too skilled	Responses in this category indicate that they are unsure since they believe hackers, companies and criminals are too skilled.	<i>"There can be other new hacking methods"</i>

Do you feel confident you can protect your smartwatch data?* *Continued*

		Perfect protection is impossible	Responses in this category indicate that they are unsure since they believe complete data protection is impossible.	<i>"A simple lock screen can only do so much"</i>
	Not previously acted		Responses in this category indicate that they are unsure since the user has never considered or tried to protect their data.	<i>"I don't actively do anything protect it"</i>
	Ignores policies		Responses in this category indicate that they are unsure since they do not always read policies, permissions and Terms & Conditions.	<i>"I don't always read the update info and permissions"</i>
	Device vulnerabilities		Responses in this category indicate that they are unsure since a device vulnerability might undermine their protection.	<i>"The information can be accessed if the manufacturer is untrustworthy"</i>
	Trust the vendor		Responses in this category indicate that they are unsure since protection is not necessary if you trust the product.	<i>"Trusting technology"</i>
	Unsure		Responses in this category indicate that they are unsure but cannot provide a reason for this opinion.	<i>"Not that sure"</i>
	Not applicable to response		Responses in this category are either too concise, too vague, unrelated to the question or not clear in their rationale.	<i>"N/a"</i>
No	Insufficiently skilled	Not informed	Responses in this category indicate that they lack confidence since they do not know enough about the topic.	<i>"Again, I have no idea how it all works"</i>
		Too complex	Responses in this category indicate that they lack confidence since protecting data is complex.	<i>"No, I feel it's a very difficult task"</i>
	Infeasible	Adversaries are too skilled	Responses in this category indicate that they lack confidence since hackers, companies and criminals are skilled in accessing data.	<i>"The companies always have a way to get your data"</i>
		Perfect protection is impossible	Responses in this category indicate that they lack confidence since they believe complete data protection is impossible.	<i>"Even with security measures, is it really safe"</i>
	Not acted previously		Responses in this category indicate that they lack confidence since the user has never considered or tried to protect their data.	<i>"I never thought about it until now"</i>
	Ignores policies		Responses in this category indicate that they lack confidence since they do not always read policies, permissions and Terms & Conditions.	<i>"I don't read ever agreement when downloading an app"</i>
	Device vulnerabilities		Responses in this category indicate that they lack confidence since a device vulnerability might undermine their protection.	<i>"I feel there are vulnerabilities I don't know about"</i>

Do you feel confident you can protect your smartwatch data?* *Continued*

Unsure	Responses in this category indicate that they lack confidence but cannot provide a reason for this opinion.	“Not sure”
Not applicable to response	Responses in this category are either too concise, too vague, unrelated to the question or not clear in their rationale.	“N/A”

Table A.26: Do you think taking action to protect your smartwatch data is worth the effort?*

Theme	Subtheme	Subsubtheme	Definition	Quote
Yes	Protection	Protect data	Responses in this category indicate that taking action is justified by the goal of protecting personal data.	“It is important that my data is protected”
		Reduce personal risk	Responses in this category indicate that taking action is justified by seeking to reduce risk.	“Yes, it can save you from having a problem that can be get really bad”
	Privacy	Values principle	Responses in this category indicate that taking action is justified by the user valuing the principle of privacy or security.	“My data is supposed to be private”
		Sensitive data	Responses in this category indicate that taking action is justified by the fact that their data is sensitive.	“Because it has some very personal information”
	Cautious		Responses in this category indicate that taking action is justified since it is prudent to protect data and be cautious.	“Better to be safe than sorry”
	Little effort		Responses in this category indicate that taking action is justified by the fact that protection requires little effort.	“It hardly takes anytime at all to protect my data”
	Helps feel secure		Responses in this category indicate that taking action is justified by the fact it helps the user feel safe or secure.	“So I feel more secure to enter my data into my watch”
	Not applicable to response		Responses in this category are either too concise, too vague, unrelated to the question or not clear in their rationale.	“Coz it is”
Unsure	Little risk	Low threat	Responses in this category indicate that they are unsure that taking action is justified since they do not face a serious threat.	“That data isn’t targeted enough”
		Data lacks sensitivity	Responses in this category indicate that they are unsure that taking action is justified since their data is not sensitive and therefore unlikely to be coveted.	“My smart watch doesn’t run any apps that collect personal data”
	Too complex		Responses in this category indicate that they are unsure that taking action is justified since they believe the matter is too complex.	“Not sure how much action is required”

Do you think taking action to protect your smartwatch data is worth the effort?* *Continued*

	Do not know risk	Responses in this category indicate that they are unsure that taking action is justified since they are not aware of the degree of risk.	<i>"Not sure what's at risk"</i>	
	Depends on data sensitivity	Responses in this category indicate that the user's willingness to act depends on how sensitive their data is.	<i>"If I keep sensitive data and passwords then I would protect it"</i>	
	Not sure why data would be desired	Responses in this category indicate that they are unsure that taking action is justified since they do not understand why their data is desirable.	<i>"Not sure how someone could use the data"</i>	
	Data accessible regardless	Responses in this category indicate that they are unsure that taking action is justified since personal data can be accessed through alternative means.	<i>"Data is sold anyway"</i>	
	Unsure	Responses in this category indicate that they are unsure that taking action is justified but do not have clear rationale.	<i>"Don't know"</i>	
	Not applicable to response	Responses in this category are either too concise, too vague, unrelated to the question or not clear in their rationale.	<i>"N/a"</i>	
No	Little risk	Low threat	Responses in this category indicate that they do not believe that taking action is justified since they do not face a serious threat.	<i>"I don't think the threat to steal that data is all that great"</i>
		Data lacks sensitivity	Responses in this category indicate that they do not believe that taking action is justified since their data is not sensitive and therefore unlikely to be coveted.	<i>"I don't have much personal information on my smart watch"</i>
	Too complex		Responses in this category indicate that they do not believe that taking action is justified since they believe the matter is too complex.	<i>"At the end of the day, there is no conceivable way to protect all of your data"</i>
	Data accessible regardless		Responses in this category indicate that they do not believe that taking action is justified since personal data can be accessed through alternative means.	<i>"Companies are getting that data no matter what"</i>
	Loss of convenience		Responses in this category indicate that they do not believe that taking action is justified since protection would impede functionality.	<i>"I lose out on a lot by turning off settings that could share my data"</i>
	Nothing to hide		Responses in this category indicate that they do not believe that taking action is justified since they believe that they have nothing to hide.	<i>"I have nothing to hide"</i>
	Trust the smartwatch vendors		Responses in this category indicate that they do not believe that taking action is justified since they trust the smartwatch companies.	<i>"I trust in the companies"</i>
	Unsure		Responses in this category indicate that they do not believe that taking action is justified but do not have clear rationale.	<i>"Not sure"</i>

Do you think taking action to protect your smartwatch data is worth the effort?* *Continued*

	Not applicable to response	Responses in this category are either too concise, too vague, unrelated to the question or not clear in their rationale.	“Anti-trust issues”
--	----------------------------	--	---------------------

Table A.27: Do you think your smartwatch data faces a realistic threat?*

Theme	Subtheme	Subsubtheme	Definition	Quote
Yes	Particular attacks	Sophisticated attacks	Responses in this category believe a threat is encountered since sophisticated hackers can access data.	“ <i>Hacking and identity thieves get more advanced every day</i> ”
		Theft	Responses in this category believe a threat is encountered since smartwatches and their data can be stolen.	“ <i>It is small and could be easily lost or stolen</i> ”
		Location tracking	Responses in this category believe a threat is encountered since parties wish to access location or track geographical positions.	“ <i>People can track your movements</i> ”
		Identity theft / fraud	Responses in this category believe a threat is encountered since parties wish to commit fraud and identity theft.	“ <i>I think that your data could be breached and used to steal your identity</i> ”
		Data selling/sharing	Responses in this category believe a threat is encountered since companies and other parties would like to sell and share data.	“ <i>Companies are willing to sell our data for profit</i> ”
	Vulnerabilities	Insecure smartwatches	Responses in this category believe a threat is encountered since smartwatches are not completely secured.	“ <i>Because it’s not really protected</i> ”
		Technology is insecure	Responses in this category believe a threat is encountered since technology in general are not completely secured.	“ <i>Everything electronic (especially connected to the internet) is hackable</i> ”
		Connected to smartphone	Responses in this category believe a threat is encountered since the smartwatch is linked to a smartphone.	“ <i>The data on my watch is synced with what my phone is carrying</i> ”
	Sensitive data		Responses in this category believe a threat is encountered since the smartwatch contains sensitive data. Therefore, it becomes a lucrative target for attackers.	“ <i>There is valuable information there which would be of interest to some people</i> ”
	Protection is never perfect		Responses in this category believe a threat is encountered since absolute protection cannot be guaranteed.	“ <i>There are always methods to circumvent security</i> ”

Do you think your smartwatch data faces a realistic threat?* *Continued*

	Unsure		Responses in this category believe a threat is encountered but cannot provide their rationale.	<i>"Not sure"</i>
	Not applicable to response		Responses in this category are either too concise, too vague, unrelated to the question or not clear in their rationale.	<i>"Very good"</i>
Unsure	Likely	Protection is never perfect	Responses in this category are unsure whether threat is encountered, but it is possible since protection is never perfect.	<i>"It's hard to know, with so many new security threats"</i>
		Potential damage	Responses in this category are unsure whether threat is encountered, but it is possible that damage can be inflicted.	<i>"Apps can be used to obtain personal information"</i>
	Not sure	Not informed	Responses in this category are unsure whether threat is encountered since they are not sufficiently informed.	<i>"I feel like I don't know enough about safety"</i>
		Unsure	Responses in this category are unsure whether threat is encountered and cannot provide their rationale.	<i>"I honestly don't know"</i>
		Not applicable to response	Responses in this category are either too concise, too vague, unrelated to the question or not clear in their rationale.	<i>"Because"</i>
	Unlikely	Data not sensitive	Responses in this category are unsure whether threat is encountered, but it is unlikely since smartwatch data is not sensitive.	<i>"Again not sure what they would get from the data"</i>
		Little perceived risk	Responses in this category are unsure whether threat is encountered, but it is unlikely since little risk is perceived.	<i>"I've not heard anything about data being stolen from a watch"</i>
		Adequately protected	Responses in this category are unsure whether threat is encountered, but it is unlikely since smartwatches are sufficiently protected.	<i>"I think smart watches have suitable security measures"</i>
		Easier/other alternatives	Responses in this category are unsure whether threat is encountered, but it is unlikely since other technologies are easier or more profitable to target.	<i>"I think phones are more a target than watches at this time"</i>
	No	Little interest	Data not sensitive	Responses in this category doubt that smartwatches are under threat since little sensitive data is stored.
Not important person			Responses in this category doubt that their smartwatch is under threat since they are not a person of importance.	<i>"Because I am nothing special"</i>
Little perceived risk		Responses in this category doubt that smartwatches are under threat since they perceive little risk.	<i>"I have never been attacked before"</i>	

Do you think your smartwatch data faces a realistic threat?* *Continued*

Adequately protected	Responses in this category doubt that smartwatches are under threat since they believe they are secure enough.	<i>"All the programs and the updates are uploaded properly"</i>
Easier/other alternatives	Responses in this category doubt that smartwatches are under threat since other devices are easier, more profitable or more accessible to target.	<i>"Other devices are attacked more"</i>
Trust smartwatch company	Responses in this category doubt that smartwatches are under threat they trust the manufacturer of their device.	<i>"The protection on this operation system is very secure"</i>
Unsure	Responses in this category doubt that smartwatches are under threat but cannot provide their rationale.	<i>"Not sure"</i>
Not applicable to response	Responses in this category are either too concise, too vague, unrelated to the question or not clear in their rationale.	<i>"Just that"</i>

Table A.28: What would influence you to take greater action to protect your smartwatch data?

Theme	Subtheme	Definition	Quote
Damage	Personally affected/targeted	Responses in this category indicate that action would be encouraged if the user was personally affected or targeted by an attack.	<i>"If I knew it was being accessed without permission"</i>
	Friends affected/targeted	Responses in this category indicate that action would be encouraged if a user's friend or family was personally affected or targeted by an attack.	<i>"If a friend or somebody I knew had their data stolen"</i>
	High-profile damage	Responses in this category indicate that action would be encouraged if a high-profile and publicised attack took place.	<i>"If there was a big scandal"</i>
Information	Threats proven/demonstrated	Responses in this category indicate that action would be encouraged if privacy threats could be clearly proven or demonstrated.	<i>"An example of what could go wrong with low protection"</i>
	More details	Responses in this category indicate that action would be encouraged if the user received additional information.	<i>"I think more information needs to be made available"</i>
	Vendor warnings	Responses in this category indicate that action would be encouraged if the user received information/warnings from their smartwatch vendor.	<i>"If there was more information that the company provided"</i>

What would influence you to take greater action to protect your smartwatch data? *Continued*

	News reports	Responses in this category indicate that action would be encouraged if the user read news reports or articles about privacy risks.	<i>“News about not ethical use of data”</i>
	Awareness advertisement	Responses in this category indicate that action would be encouraged if the user received additional information from awareness campaigns/initiatives.	<i>“Information campaigns”</i>
	Friends advice	Responses in this category indicate that action would be encouraged if a user’s friend advised them to take more control over their privacy.	<i>“Advice from friends”</i>
	New helpful features	Responses in this category indicate that action would be encouraged by the addition of new features or settings.	<i>“If there were new or additional security measures”</i>
	Increased risk	Responses in this category indicate that action would be encouraged in response to an increased degree of privacy risk.	<i>“If I was more at risk of having it invaded or stolen”</i>
	Storing sensitive data	Responses in this category indicate that action would be encouraged if the user had sensitive data on their watch.	<i>“If I had sensitive information on my smartwatch”</i>
	Advantages/incentives	Responses in this category indicate that action would be encouraged if privacy settings offered additional advantages or incentives.	<i>“Rewards for having security setup”</i>
	Nothing	Responses in this category indicate that nothing or little could be done to encourage greater action.	<i>“Nothing really”</i>
	Already take action	Responses in this category indicate that they have used tools to protect their data.	<i>“I’ve done everything I can now”</i>
	Principle	Responses in this category indicate that they are encouraged by the principle of security, privacy or data protection.	<i>“Importance of privacy”</i>
	Unsure	Responses in this category are unsure what could encourage them to take greater action.	<i>“I don’t know”</i>
	Not applicable to response	Responses in this category either do not relate to the question, are too vague, excessively concise, or do not provide rationale.	<i>“N/a”</i>

A.4 Chapter 8: Longitudinal Study

Concern questionnaire tables

Table A.29: Stranger Access Scenario Pair - Data Access

Theme	Subtheme	Definition	Quote
Unconcerned	Aid watch return	Such responses believe that stranger access could be beneficial, as they could assist in returning the watch.	<i>“But may help it to be returned if it had been lost”</i>
	Screen lock protection	Such responses believe damage from access is limited since a screen lock (PIN, password or pattern) is enabled on the watch.	<i>“And I have a lock on it”</i>
	Lack of sensitive data	Such responses are unconcerned since they believe strangers could not extract much value from their smartwatch data.	<i>“Not sure they’d get much out of it”</i>
Contingent	Depends on sensitivity	Such responses state their concern depends on the sensitivity of the accessed data. Individuals might not mind certain details being read, but would oppose other data being accessed.	<i>“If they had my exercise information I wouldn’t mind much”</i>
	Depends on accessor	Such responses state their concern depends on the abilities of the stranger accessing data.	<i>“Depends on the stranger! Someone without computing skills - unconcerned as I enabled a PIN”</i>
Limited damage	Will not sell unlike companies	Such responses believe damage from access is limited since ordinary individuals won’t sell data. This is unlike tech companies.	<i>“The average stranger wouldn’t sell this data to organisations”</i>
	Limited smartwatch functionality	Such responses believe damage from access is limited since smartwatches only have a small array of functionality.	<i>“I would mind less, the smartwatch is much more restricted”</i>
	Remotely disconnect accounts	Such responses believe damage from access is limited since app accounts can be disconnected remotely.	<i>“I would think I would be able to de-link my account from the device remotely”</i>
	Hardware more valuable than data	Such responses believe damage from access is limited since strangers are more likely to profit from the hardware rather than the data.	<i>“I’d hope they would be looking to sell the watch, not steal my data”</i>
Greater damage	Principle	Such responses are concerned because of the ethical issues of a stranger accessing their data. They believe such an event is invasive and a breach of their right to privacy.	<i>“Breach of personal privacy”</i>
	Fraud	Such responses are concerned because a stranger could use the smartwatch data to commit fraud, such as identity theft.	<i>“There are many kinds of fraud easily doable just knowing few personal data”</i>

Stranger Access Scenario Pair - Data Access *Continued*

	Security risk	Such responses are concerned due to security or safety risks. Individuals may believe they could be harmed as a result of a stranger reading this data.	<i>“They may use it to harm me or extortion me”</i>
	Personal data	Such responses are concerned because they believe their personal data is private and sensitive.	<i>“I don’t want a stranger know...my calendar (or have sensitive data)”</i>
	Identification	Such responses are concerned because individuals might be identified based on their data.	<i>“I could be easily identified”</i>
	Inconvenience	Such responses are concerned because strangers accessing their data might cause inconvenience.	<i>“Also, if I use Android program, it might be inconvenient”</i>

Table A.30: Stranger Access Scenario Pair - App Usage

Theme	Subtheme	Definition	Quote
Unconcerned	Lack of sensitive data	Such responses are not concerned as they believe their apps don’t provide access to sensitive or serious data.	<i>“I don’t have any serious apps there”</i>
Contingent	Depends on app	Such responses state their concerns depend on which app is used by a stranger. If a remark qualifies concern by referring to particular app, this category is relevant.	<i>“But the answer depends on the app itself”</i>
	Depends on damage	Such responses state their concerns depend on the damage inflicted by app usage. Issues could include fraudulent emails, financial harm or risk to a family.	<i>“If the stranger could send emails using my account, I will be concerned”</i>
Limited damage	Screen lock protection	Such responses believe damage from app use is limited since a screen lock (PIN, password or pattern) is enabled on the watch.	<i>“PIN feature helps protect against this”</i>
	Remotely disconnect accounts	Such responses believe damage from app use is limited since app accounts can be disconnected remotely. This could be achieved by changing account passwords.	<i>“I think this one is easily solvable by changing account passwords”</i>
Greater damage	Principle	Such responses are concerned because of the ethical issues of a stranger using their applications. They believe such an action is invasive and a breach of their right to privacy.	<i>“Bit creepy!”</i>
	Identity theft	Such responses are concerned because a stranger could commit identity theft through using a victim’s data.	<i>“Identity theft is my worst fear”</i>

Stranger Access Scenario Pair - App Usage *Continued*

Impersonation	Such responses are concerned because a stranger could impersonate a victim. Then the stranger could undertake actions on the victim's behalf and trick their colleagues.	<i>"My worst fear is being impersonated"</i>
Access personal data	Such responses are concerned about their personal data being accessed by a stranger.	<i>"Could also access my personal data"</i>
Security risk	Such responses are concerned due to security or safety risks. Individuals may be worried that accessed data might cause financial damage.	<i>"I have links to messages, social media and could have Android Pay enabled - could put me at risk"</i>
Ruin fitness statistics	Such responses are concerned because their fitness statistics might be adjusted by a stranger.	<i>"Would ruin my running stats"</i>
Could spend money	Such responses are concerned since they believe a stranger could spend or steal their money through smartwatch apps.	<i>"They could download more apps from the store that might enable them to spend my money"</i>
Extortion	Such responses are concerned since they believe a stranger could use personal data to extort, incriminate, bribe or blackmail a victim.	<i>"They could download more apps from the store that might enable them to...incriminate me"</i>

Table A.31: Location Tracking Scenario Pair - Location Tracking

Theme	Subtheme	Definition	Quote
Unconcerned	Provides benefits	Responses in this category believe that location tracking can provide benefits, such as useful functionality.	<i>"Would be useful to track my running"</i>
	Benign usage	Responses in this category believe that location tracking would not necessarily be detrimental. Individuals believe that even if this data is collected, it will not be harmful.	<i>"I have no reason to think that this data would be used in any way to harm me"</i>
	Required for functionality	Responses in this category believe that location tracking is necessary to provide certain services. If tracking was not undertaken, the app would not work.	<i>"Wouldn't bother me too much as many apps need location to function properly"</i>
	Can be disabled	Responses in this category are not concerned, since they believe that location tracking can be simply disabled.	<i>"Can turn Wi-Fi/GPS off if I really didn't want them to know where I am"</i>
	Can access data regardless	Responses in this category are not concerned, since they believe that location data is already collected by many other devices.	<i>"Location data is anyway tracked in the modern day world"</i>

Location Tracking Scenario Pair - Location Tracking *Continued*

Contingent	Depends on use	Responses in this category state their concern depends on how location data is used. While benign processes might be accepted, invasive practices would be opposed.	<i>“However, could be more concerned based on how they used my location data”</i>
	Depends on the necessity	Responses in this category state their concern depends on the necessity of the tracking. If an app requires this data to provide a service, that is more acceptable.	<i>“If it is a GPS app for sports, that is needed to provide the best user experience”</i>
	Depends if optional	Responses in this category state their concerns depend on whether tracking is optional. If the location monitoring can be turned on or off at the user’s command, then the practice might be more acceptable.	<i>“I want to have the option for this not to happen”</i>
	Depends if informed	Responses in this category state their concerns depend on whether users are aware they are being monitored. If individuals know that their location is tracked, then they can make decisions accordingly.	<i>“If I know about...then it’s ok”</i>
	Depends if beneficial	Responses in this category state their concerns depend on whether they receive any benefits from an app accessing their location.	<i>“If there is an advantage for me I don’t care”</i>
Concerned	Surveillance	Responses in this category are concerned about their movements being tracked.	<i>“Not too keen on people keeping track of my movements”</i>
	Leak risk	Responses in this category are concerned because tracked data might be accessed by hackers. If a wide range of apps can monitor location, there are more entities that could leak those details.	<i>“There might be a hack which may cause a breach of both my personal security and privacy”</i>
	Data selling	Responses in this category are concerned about their location data being sold to others.	<i>“Would be concerned that they may sell access to that data to other organisations”</i>
	Non-location tracking	Responses in this category are concerned since location tracking might imply that other details are also being monitored.	<i>“I would be worried they may be tracking something else!”</i>
	Personal use	Responses in this category would prefer location tracking to be reserved to personal use. In this manner, benefits would be exclusive to the individual not to companies.	<i>“I would prefer this being for my own personal use only”</i>
	Principle	Such responses are concerned because of the ethical issues of privacy invasion. These individuals might believe their location is private and that others have no right to access it.	<i>“I would prefer that they don’t know for privacy reasons”</i>

Location Tracking Scenario Pair - Location Tracking *Continued*

	Targeted advertising	Such responses are concerned that collected location data might be used for targeted advertising purposes.	<i>“They could target some information upon where I’ve been or where I’ve shopped”</i>
	Identification	Such responses are concerned because individuals might be identified based on their location data.	<i>“Can be used to identify me”</i>
	No benefit to user	Such responses are concerned since apps receive benefits from location data, rather than the user.	<i>“Only benefits them, not me”</i>

Table A.32: Location Tracking Scenario Pair - Location Sharing

Theme	Subtheme	Definition	Quote
Unconcerned	Benign usage	Responses in this category do not believe that such an action would necessarily lead to harmful consequences.	<i>“I doubt harmful”</i>
	Accepted reality	Responses in this category accept that location sharing occurs as a reality. If companies have access to data, they will try to profit from it.	<i>“Location tracking is a fact of life in our modern technological world”</i>
	Use GPS rarely	Responses in this category are unconcerned about location sharing since the participant uses GPS rarely.	<i>“As I only use this feature occasionally it wouldn’t bother me too much”</i>
Contingent	Depends if anonymous	Such responses state their concerns depend on whether the shared location data is anonymised.	<i>“I would prefer it to be anonymised, so I couldn’t be identified personally”</i>
	Depends if beneficial	Such responses state their concerns depend on whether they receive any benefits from an app company sharing their location.	<i>“It depends on what I get back from it”</i>
	Depends if consented	Such responses state their concerns depend on whether they have agreed that their location data may be shared.	<i>“Need to get me to agree to use this info”</i>
	Depends on granularity	Such responses state their concerns depend on how vague or granular the shared data is.	<i>“Vague movements are ok”</i>
Concerned	Principle	Such responses are concerned because of the ethical issues of privacy invasion. These individuals might believe their location is private and that others have no right to access it.	<i>“Breach of personal privacy”</i>
	Illegal	Such responses are concerned over the perceived illegality of sharing location data.	<i>“It also doesn’t seem terribly legal”</i>
	Surveillance	Responses in this category are concerned about their movements being tracked.	<i>“I don’t want them to know where I am at any one time”</i>

Location Tracking Scenario Pair - Location Sharing *Continued*

Targeted advertising	Such responses are concerned that shared location data might be used for targeted advertising purposes.	<i>“Because they could target products”</i>
Personal data	Such responses are concerned because their location data is personal and sensitive. These individuals may be worried about this private information being held by others.	<i>“My location is a personal information”</i>
Security risk	Such responses are concerned due to security or safety risks. Individuals might believe location sharing is worrying or practically dangerous.	<i>“Unauthorised sharing of my precise location data worries me on...practical grounds”</i>
Uninformed	Such responses are concerned due to a lack of knowledge. Individuals might not be informed who is accessing their data or why.	<i>“I don’t know who has this data or why”</i>
Data can be deanonymised	Such responses are concerned because they believe that shared location data could be deanonymised.	<i>“Easy to deanonymise”</i>

Table A.33: App Access Scenario Pair - Data Access

Theme	Subtheme	Definition	Quote
Unconcerned	Can access data regardless	Such responses are not concerned since they assume similar data can already be accessed through other technologies.	<i>“They already can with phone, or at least I assume so”</i>
	Access required for functionality	Responses in this category believe that some personal data is necessary to provide certain services.	<i>“While certain personal data may be necessary”</i>
	Lack of sensitive data	Such responses are not concerned as they believe their watches don’t possess sensitive data.	<i>“There’s nothing on my smartwatch that is particularly sensitive”</i>
Contingent	Depends if optional	Responses in this category state their concerns depend on whether the access to data is optional.	<i>“Would need to be opt-in only”</i>
	Depends if informed	Responses in this category state their concerns depend on whether users are aware of what companies are doing.	<i>“However, if I know precisely what they’re doing...then OK!”</i>
	Depends if aggregated	Such responses state their concerns depend on whether their data is aggregated with other records. If so, their individual details are not as identifiable.	<i>“I wouldn’t mind so much if used...at the aggregate level”</i>

App Access Scenario Pair - Data Access *Continued*

	Depends if linked to location	Such responses state their concerns depend on whether personal data is connected to their location. If so, it might be more sensitive and damaging.	<i>“Personal data ie heartrate may be more concerning if linked with location”</i>
	Depends if used responsibly	Such responses state their concerns depend on whether the accessed data is used in a responsible manner.	<i>“I wouldn’t mind so much if used responsibly”</i>
	Depends if sensitive data	Such responses state their concern depends on the sensitivity of the accessed data. Individuals might not mind certain details being read, but would oppose other data being accessed.	<i>“It depends what form of personal data”</i>
	Depends if beneficial	Such responses state their concerns depend on whether they receive any advantages from an app accessing their data.	<i>“If there is an advantage for me I don’t care”</i>
Concerned	Principle	Such responses are concerned because of the ethical issues of their data being accessed. They value their data and believe such an action is invasive and a breach of their privacy.	<i>“I would want them to respect my privacy”</i>
	Targeted advertising	Such responses are concerned that app data might be used for targeted advertising purposes.	<i>“Could be used to target ads to me”</i>
	Leak risk	Such responses are concerned that collected data might be hacked, therefore leaking sensitive details.	<i>“Also, I am not sure it’s protected well”</i>
	Personal data	Such responses are concerned because they believe their personal data is private and personal.	<i>“This should remain private as it is personal data”</i>
	No benefit	Such responses are concerned since they believe the data collection will have no benefit to the user.	<i>“I can’t really imagine what they would use this for that is of benefit to me”</i>
	Data selling	Responses in this category are concerned about their data being sold to others.	<i>“Wouldn’t trust the companies to not sell the information”</i>
	Unconsented	Such responses are concerned due to a perceived lack of consent for data access.	<i>“I would not want app companies that I do not allow to have access to my personal data”</i>
	Identification	Such responses are concerned because individuals might be identified based on their data.	<i>“Can be used to identify me”</i>

App Access Scenario Pair - Data Access *Continued*

	Illegal	Such responses are concerned over the perceived illegality of accessing watch data.	<i>“Also, I’m pretty sure that’s a crime”</i>
--	---------	---	---

Table A.34: App Access Scenario Pair - Data Sharing

Theme	Subtheme	Definition	Quote
Unconcerned	Accepted as reality	Such responses accept data sharing as being a modern reality, whether fortunate or unfortunate. If companies can profit from data, they will.	<i>“Although I am sure they largely can and do”</i>
	Technology issue in general	Such responses believe that data sharing issues are not specific to smartwatches, but modern technology in general.	<i>“Having any electronic device or internet presence makes this a reality”</i>
	Allows functionality	Such responses are unconcerned since data sharing is perceived to deliver quality services.	<i>“Normally the data shared is used to allow companies to analyse and then provide higher quality”</i>
	Mostly for advertisements	Such responses are unconcerned since data sharing is mainly used for marketing purposes, which is not considered malicious.	<i>“I think it is mostly for marketing purposes which is not too malicious”</i>
Contingent	Depends on the data	Such responses state their concerns depend on what data (and the sensitivity of that data) is being shared with other companies.	<i>“Bit concerned, but it depends exactly what data”</i>
	Depends on the recipient	Such responses state their concerns depend on the identity of the entity receiving the shared data.	<i>“Bit concerned, but it depends...the companies they share it with”</i>
	Depends if aggregated	Such responses state their concerns depend on whether their data is aggregated with other records. If so, their individual details are not as identifiable.	<i>“I wouldn’t mind if it was aggregated”</i>
	Depends if anonymised	Such responses state their concerns depend on whether the shared data is anonymised.	<i>“I would prefer it to be anonymised, so I couldn’t be identified personally”</i>
	Depends if beneficial	Such responses state their concerns depend on whether they get any benefits back from their data being shared.	<i>“Again, it depends on what I get back from it”</i>
Concerned	Principle	Such responses are concerned because of the ethical issues of their data being shared. They believe such an action is invasive and a breach of their right to privacy.	<i>“Breach of privacy”</i>
	Unconsented	Such responses are concerned since they don’t believe they agreed to data sharing.	<i>“Probably didn’t agree to such sharing in the first instance”</i>

App Access Scenario Pair - Data Sharing *Continued*

Lack of control	Such responses are concerned due to a perceived lack of control over their data. They might not get a choice over who can access their information.	<i>"I don't get to choose which companies"</i>
Targeted advertising	Such responses are concerned that shared data might be used for targeted advertising purposes.	<i>"However, it may lead to annoying advertisements potentially"</i>
Surveillance	Responses in this category are concerned about being tracked based on their shared data.	<i>"I would be worried that wider organisations would know where I lived"</i>
Leak risk	Such responses are concerned that collected data might be stolen or misused, therefore leaking details.	<i>"I don't necessarily trust the way they'd store my data"</i>
Unnecessary	Such responses are concerned since they believe that data sharing is not necessary.	<i>"No need"</i>
Personal data	Such responses are concerned because they believe their personal data should remain private.	<i>"I do not want my personal data to be running free"</i>

Game evaluation tables

Table A.35: What things did you like about the smartwatch game?

Theme	Subtheme	Definition	Quote
Educational	General	This category concerns educational merit in general, rather than referring to a particular topic. An individual would gain greater knowledge of their smartwatch as a whole.	<i>"It made me learn how to use the watch and showed me features I wouldn't have otherwise used"</i>
	Security/privacy	This category concerns education in the context of security or privacy. The lessons learned might help a participant to protect their data.	<i>"I liked that it gave me a better understanding of how to protect myself when using it"</i>
Usability	Interaction	This category concerns usability in terms of user interaction. The app would be easy and straightforward to use, rather than performant, simple or smooth.	<i>"It was interactive and straightforward to play"</i>
	Simplicity	This category concerns the ease of understanding the game. It does not relate to the user experience, but how simply the app can be appreciated.	<i>"Easy to understand"</i>
	Performance	This category concerns how smoothly and correctly the app functions.	<i>"It never crashed and went smoothly"</i>
Design	Aesthetics	This category concerns how attractive the app appears.	<i>"Cute interface"</i>
	Layout	This category relates to the way in which the application has been laid out and designed.	<i>"The fake settings menu"</i>
	Character	This category concerns the characters in the game; both NPC and player-controlled.	<i>"Element of personalisation (character)"</i>
	Flexibility	This category concerns how easily the game could be applied to a different context.	<i>"Could be used for different games (for different settings like factory, etc.)"</i>

Table A.36: What things did you not like about the smartwatch game?

Theme	Subtheme	Definition	Quote
Boring	Theme	This category concerns the general topic of the game, that of shopping, being boring.	<i>"Setting (shopping isn't that cool)"</i>
	Repetitive	This category concerns elements of the game being overly repetitive.	<i>"It got a bit repetitive after a while"</i>
	Challenge variety	This category concerns a lack of diversity and variety in game challenges.	<i>"Very little variety of challenges"</i>
	Map variety	This category concerns a lack of diversity and variety in map layout.	<i>"Morning/Afternoon/Night seemed to be exactly the same course"</i>

What things did you not like about the smartwatch game? *Continued*

Usability issues	Navigation	This category concerns navigation and movement issues within the game.	<i>“Not very intuitive in terms of movement”</i>
	Screen lock challenges	This category concerns usability issues encountered when trying to complete the PIN or password challenges.	<i>“Problem with the PIN/password setting challenges”</i>
Difficulty	Too easy	This category concerns complaints that elements of the game were excessively easy to complete.	<i>“Questions from citizens were really easy”</i>
	Too difficult	This category concerns complaints that elements of the game were excessively difficult to complete.	<i>“Found the challenges get too hard on Afternoon setting”</i>

Table A.37: What about the game would you like to see improved?

Theme	Subtheme	Definition	Quote
Extend game	Challenges	This category concerns suggestions to include additional challenges, or to diversify their range.	<i>“Possibly create more challenges to choose from”</i>
	Questions	This category concerns suggestions to include additional questions, or to diversify their range.	<i>“There could be...a greater variety in questions”</i>
	Levels	This category concerns suggestions to include additional levels, or to diversify their range.	<i>“There could be more levels”</i>
	Difficulties	This category concerns suggestions to include additional difficulties.	<i>“Perhaps some...extra difficulties to give it a longer play time”</i>
Improve usability	Fix challenge bugs	This category concerns suggestions to fix the user interface errors within the challenges/tasks.	<i>“Sensitivity of the ‘settings’ page especially the PIN/password input”</i>
	Enhance interface	This category concerns suggestions to further enhance the usability of the user interface.	<i>“Menu scrolling”</i>
	Reduce challenge difficulty	This category concerns suggestions to increase the amount of allotted time to the longer challenges.	<i>“More time for the ‘longer’ challenges”</i>
Add new features	Battle system	This category concerns suggestions to introduce a ‘battle system’ into the game.	<i>“Battle system”</i>
	Greater maze	This category concerns suggestions to make the map layout more of an advanced maze.	<i>“More of a maze”</i>
	Level menu	This category concerns suggestions to add an option to allow players to start at higher level when they next play the game.	<i>“Option to start at a harder level each time”</i>

Posttest interview tables

Table A.38: Why did you choose to participate in the study?

Theme	Subtheme	Definition	Quote
Opportunity to play with watch	Curiosity	Responses in this category participated in this study since they were curious, interested and intrigued to try out a smartwatch. This category does not discuss a temptation to buy the device, or that they are checking whether it matches their needs.	<i>“Curiosity, yeah. Wanted to see what you could do with a smartwatch”</i>
	Tempted to buy watch before	Responses in this category participated in this study since they were previously tempted or thinking about purchasing a smartwatch. This concerns previous thoughts of making a purchase, rather than checking suitability.	<i>“I always wanted to buy a smartwatch”</i>
	Check suitability/value	Responses in this category participated in this study since they wanted to test whether a smartwatch was suitable for them or good value.	<i>“Whether it was something that I would consider buying”</i>
	Approve of watch manufacturer	Responses in this category participated in this study since they were a fan of the manufacturer (Huawei) of the smartwatches.	<i>“I think Huawei is a good company”</i>
Incentive of compensation		Responses in this category participated in this study since they were incentivised by receiving Amazon vouchers as compensation.	<i>“The Amazon vouchers were pretty good”</i>
Participate in academic research		Responses in this category participated in this study since they like to take part in academic studies.	<i>“I think it’s a nice way to give back to the academic way”</i>

Table A.39: Do you feel the background StudyService app affected your behaviour?

Theme	Definition	Quote
No	Responses in this category think that the StudyService did not affect their behaviour.	<i>“Not at all, no. It was just sitting there”</i>
Checked the app ran	Responses in this category indicate that the participant occasionally checked the StudyService to ensure it was running correctly.	<i>“I just let it live normally and left it and just checked that it was okay”</i>

Table A.40: Do you feel you learned anything new as the study progressed?

Theme	Subtheme	Definition	Quote
Privacy features	Location privacy	Responses in this category indicated that they learned about location privacy.	<i>“I definitely feel like I’ve learned more about privacy and location”</i>
	Privacy permissions	Responses in this category indicated that they learned about the app privacy permissions on the smartwatch. This includes descriptions of how applications can or cannot access certain pieces of data.	<i>“Interesting to get to know some of the permissions features of the watch”</i>
	Screen locks	Responses in this category indicated that they learned about screen locks (either PINs, passwords or patterns) on the smartwatch.	<i>“I think after that I put on a PIN because balancing convenience versus safety”</i>
General smartwatch	Smartwatch/settings use	Responses in this category indicated that they learned about how to use and navigate general functions on the smartwatch. This might include the Android operating system or the watch settings.	<i>“I picked up actually how to use different functions and so on”</i>
	Smartwatch apps	Responses in this category indicated explicitly that they learned about the variety of smartwatch apps.	<i>“Yeah, mostly the apps on the watch”</i>
	Charging	Responses in this category indicated that they learned about how to charge the smartwatch.	<i>“Charging”</i>
	Voice-to-text feature	Responses in this category indicated that they learned about how to convert speech to text and send a message.	<i>“I learned that you could speak into a watch and send a message”</i>
	Synching watch and phone	Responses in this category indicated that they learned about synching data between the smartwatch and their smartphone.	<i>“It was quite interesting seeing how literally just connecting it via Bluetooth via my phone”</i>
Control group settings	Smart Reply	Responses in this category indicated that they learned about the Smart Reply feature, which was mentioned in the control group game.	<i>“So I was definitely aware of where things like Activate Smart Reply”</i>
	Gestures	Responses in this category indicated that they learned about smartwatch gestures, which were mentioned in the control group game.	<i>“So I think I didn’t know you could do gestures”</i>

Do you feel you learned anything new as the study progressed? *Continued*

Their heart rate	Responses in this category indicated that they learned about their own heart rate.	<i>“And I learned what my heartrate does over time”</i>
Education from game	Responses in this category indicated that they played the game (whether it was the treatment-group version or the control-group version). This might be suggested by people claiming they didn't know certain facts in the past, but they think about them now.	<i>“The biggest learning points were when I was playing the game”</i>

Table A.41: How do you think your smartwatch's data could be accessed by others?

Theme	Subtheme	Subsubtheme	Definition	Quote
Approaches	Connections	GPS	Responses in this category believe data could be accessed via the built-in GPS services.	<i>“Well, it has its own GPS built-in”</i>
		Bluetooth	Responses in this category believe data could be accessed via Bluetooth.	<i>“It's conceivable that somebody with a Bluetooth receiver or transmitter could probably sponge data off of it”</i>
		Internet	Responses in this category believe data could be accessed when connecting to the Internet, whether WiFi or mobile (3G/4G).	<i>“I think every time you're connected to Internet I'm sure they have some kind of access to your data”</i>
	Apps	Through normal apps	Responses in this category believe data could be accessed by legitimate smartwatch apps.	<i>“Through some app that you allow them to track your location or your data”</i>
		Fake/fraudulent apps	Responses in this category believe data could be accessed through fake or fraudulent apps.	<i>“I mean, they could obviously fake an app to suggest they're doing something else”</i>
	Google account	Responses in this category believe data could be accessed through their Google account. A Google account is required to use Wear OS, and it must be connected to the watch to download apps.	<i>“Because it's connected to my Google account and so on, it could be...sold or used to advertisers”</i>	
	Software updates	Responses in this category believe data could be accessed via software updates, whether app or operating system.	<i>“Or maybe software updates maybe”</i>	

How do you think your smartwatch's data could be accessed by others? *Continued*

	Sophisticated techniques	Responses in this category believe data could be accessed through the specialised techniques of a smart adversary.	<i>"With a very smart guy behind a computer you can do whatever you want"</i>
	Not concerned	Responses in this category are not concerned about their data being accessed, since they do not believe their risk to be high.	<i>"It's not something that concerns me a lot"</i>

Table A.42: If you wanted to, what could you do to protect your smartwatch's data?

Theme	Subtheme	Subsubtheme	Definition	Quote
Approaches	Disable GPS		Responses in this category would protect their smartwatch's data by disabling, deactivating or turning off their GPS/location.	<i>"I'd probably try and deactivate the GPS tracking"</i>
	Change privacy permissions		Responses in this category would protect their smartwatch's data by reviewing and changing their app permission settings. Participants might not refer to 'permissions' explicitly, but mention app-based settings.	<i>"Well I could revoke permissions on certain apps"</i>
	Enable screen lock	PIN	Responses in this category would protect their smartwatch's data by enabling a PIN code screen lock.	<i>"Could go and set a..PIN on it again"</i>
		Password	Responses in this category would protect their smartwatch's data by enabling a password screen lock.	<i>"Could go and set a password"</i>
		Pattern	Responses in this category would protect their smartwatch's data by enabling a pattern screen lock.	<i>"Putting a pattern to unlock it"</i>
	Uninstall apps		Responses in this category would protect their smartwatch's data by uninstalling third-party apps from the device.	<i>"Probably take the other third-party apps off the phone, off the watch, sorry"</i>
	Minimise sensitive activities		Responses in this category would protect their smartwatch's data by limiting the sensitive apps and activities that they undertake on the device.	<i>"If I only limit myself to the functions that I currently use, I just won't do anything"</i>
	Reduce notifications		Responses in this category would protect their smartwatch's data by turning off or limiting the visibility of their device's notifications.	<i>"I would turn off the...notifications for example"</i>

If you wanted to, what could you do to protect your smartwatch's data? *Continued*

Log out of apps	Responses in this category would protect their smartwatch's data by logging out of important apps, since as email or messages.	<i>"I think the most important thing I'd like to log out would be my email maybe"</i>
Factory reset the watch	Responses in this category would protect their smartwatch's data by factory setting the device, removing all the information.	<i>"I think I would just do a factory reset, that would be the best solution as everything is wiped"</i>
Disable Bluetooth	Responses in this category would protect their smartwatch's data by disabling Bluetooth connectivity.	<i>"Perhaps Bluetooth connectivity, I would turn that off as well"</i>
Secure paired smartphone	Responses in this category would protect their smartwatch's data by protecting their paired smartphone.	<i>"I think it's because I assume it only work...I needed it to be connected to the phone"</i>
Change brightness	Responses in this category would protect their smartwatch's data by changing screen brightness.	<i>"Yeah, brightness"</i>
Change alarm volume	Responses in this category would protect their smartwatch's data by changing alarm volume.	<i>"Volume, alarm volume"</i>
Unsure	Responses in this category are not exactly sure how they could protect their smartwatch's data. They might not have tried protection before, or failed to check their settings.	<i>"I didn't even check that on the watch to be honest"</i>
Not concerned	Responses in this category did not consider their data being accessed, and do not care.	<i>"Don't see, don't care"</i>

Table A.43: If you wanted to limit watch access in case of loss/theft, what could you do?

Theme	Subtheme	Subsubtheme	Definition	Quote
Approaches	Enable screen lock	PIN	Responses in this category would limit access in case of loss/theft by enabling a screen lock PIN.	<i>"You could enable a PIN or a lock screen"</i>
		Password	Responses in this category would limit access in case of loss/theft by enabling a screen lock password.	<i>"Potentially some form of password"</i>
		Pattern	Responses in this category would limit access in case of loss/theft by enabling a screen lock pattern.	<i>"Yeah, I set up a screen pattern which is useful"</i>
	Track watch remotely		Responses in this category would limit access in case of loss/theft by using a watch tracker. In this case, the device could be found if it was missing or stolen.	<i>"There's probably some reverse tracker you can get"</i>
	Reset watch remotely		Responses in this category would limit access in case of loss/theft by resetting/bricking the watch remotely.	<i>"I think there was a feature that allows the resetting it to the default setting"</i>
	Change credentials remotely		Responses in this category would limit access in case of loss/theft by deactivating or changing their credentials. Therefore, functions on the watch would then stop working.	<i>"Perhaps there's some way you can deactivate or invalidate your credentials"</i>
	Disconnect through proximity		Responses in this category would limit access in case of loss/theft by simply losing proximity to the watch. Since the device is paired with a smartphone, some features will disappear as the watch leaves that region.	<i>"You need to be some proximity to your phone to use your smartwatch"</i>
	Unsure		Responses in this category are not sure how they would limit access in case of loss or theft.	<i>"I don't know"</i>

Table A.44: If you wanted to prevent apps from tracking your location, what could you do?

Theme	Subtheme	Definition	Quote
Approaches	Disable GPS/location on smartwatch	Responses in this category would prevent location tracking by disabling or turning off smartwatch GPS/location. Since the question is in the context of smartwatches, a general mention of GPS should relate to this device.	<i>"Yeah, I'd just disable the GPS"</i>
	Disable GPS on paired smartphone	Responses in this category would prevent location tracking by disabling smartphone GPS.	<i>"Well then on my phone I would have to turn off the location feature, so that's good enough"</i>
	Revoke app permissions	Responses in this category would prevent location tracking by changing/disabling the permissions of tracking apps. The description for this might relate to revoking their agreement to allow access.	<i>"And then for specific apps you could revoke their permissions"</i>
	Disconnect watch from smartphone	Responses in this category would prevent location tracking by disconnecting the watch from the paired smartphone.	<i>"Also I can disconnect the watch from my telephone"</i>
	Uninstall apps	Responses in this category would prevent location tracking by uninstalling/removing apps which monitor location.	<i>"Remove the app, if it's not vital to the functioning"</i>
	Unsure	Responses in this category are not sure of how to prevent location tracking. This might be because they lack confidence or since they have not tried before.	<i>"I haven't checked location settings, I'm not so sure actually"</i>

Table A.45: If you wanted to stop apps from reading your personal data, what could you do?

Theme	Subtheme	Definition	Quote
Approaches	Revoke app permissions	Responses in this category would prevent apps accessing data by changing/disabling the permissions for those apps.	<i>"You could revoke permissions for your contacts or your personal data, etcetera"</i>
	Uninstall apps	Responses in this category would prevent apps accessing data by uninstalling apps which they believe to place their data at risk.	<i>"Or I could uninstall the app if I felt that was needed"</i>
	Avoid/remove sensitive data on watch	Responses in this category would prevent apps accessing data by ensuring they have no sensitive data on their watch.	<i>"Ensuring I don't have any sensitive data on there"</i>

If you wanted to stop apps from reading your personal data, what could you do? *Continued*

	Disconnect watch from smartphone	Responses in this category would prevent apps accessing data by disconnecting their watch from the paired smartphone. This would stop the phone from receiving watch data.	<i>“So it’s not connected to my mobile phone”</i>
	Unsure	Responses in this category are not certain on how to prevent apps accessing data.	<i>“Not sure how to go about doing that”</i>
	Impossible	Responses in this category doubt it is possible to delete the data that apps have access to.	<i>“I also don’t think you could make it do that”</i>

Table A.46: Could you please show me how to enable a screen lock on your smartwatch?

Theme	Subtheme	Definition	Quote
Route	Direct through Settings and Personalisation	Responses in this category enabled a screen lock through a direct route, going via the ‘Settings’ menu and then the ‘Personalisation’ menu. This route may be preceded by unlocking the watch if a screen lock is enabled. This route may be extended by describing the PIN/Password/Pattern options.	<i>“I go to Settings, I scroll down to Personalisation, and then I could click Screen Lock”</i>
	Indirect route	Responses in this category enabled a screen lock through an indirect route. They succeeded in the task, but did not go straight through the ‘Settings’ menu and then the ‘Personalisation’ menu.	<i>“So I think that would be under Personalisation or Accessibility again. I’ll try Accessibility, nope”</i>
Confidence	Certain	Responses in this category were reasonably certain in the process/route of enabling a screen lock. They did not display doubts over the path they were taking. Even if technical issues were encountered, they might be certain over the route they are attempting.	<i>“Settings again, going to scroll down to Personalisation, and then Screen Lock”</i>
	Uncertain	Responses in this category had less certainty in the process/route of enabling a screen lock. They either doubted the path or guessed at certain stages.	<i>“I don’t see anything in these things so it’s not in Sound, I guess it’s Personalisation or Accessibility”</i>

Could you please show me how to enable a screen lock on your smartwatch? *Continued*

Had not enabled a screen lock before	Responses in this category admitted to having never enabled a screen lock before.	<i>“Another thing I haven’t done”</i>
--------------------------------------	---	---------------------------------------

Table A.47: Could you please show me how to disable GPS on your smartwatch?

Theme	Subtheme	Definition	Quote
Route	Direct through Settings and Connectivity	Responses in this category disabled GPS through following a direct route, via the ‘Settings’ menu and then the ‘Connectivity’ menu. This route may be preceded by unlocking the watch if a screen lock is enabled. The route may be extended by the participant then mentioning Location.	<i>“Clicking on Settings, disable GPS you said? Go down to Connectivity, Location, off”</i>
	Indirect route	Responses in this category disabled GPS through following an indirect route, potentially since they made mistakes. They completed the task but did not go straight through the ‘Settings’ menu and then the ‘Connectivity’ menu.	<i>“And I see what, maybe System? I don’t see privacy thing”</i>
Confidence	Certain	Responses in this category were reasonably certain over the process of disabling GPS. They did not display doubts over the path they were taking.	<i>“So going into Settings, and then I’m going to Connectivity, scroll down to Location and turn Location off”</i>
	Uncertain	Responses in this category were less certain over the process/route to disable GPS. They either doubted the path or guessed at certain stages.	<i>“Disable GPS? Access the menu. And probably I suspect go to Settings. Hmm”</i>
Had not disabled GPS before		Responses in this category admitted that they had not disabled GPS before.	<i>“I haven’t done this before”</i>

Table A.48: Could you please show me how to change the permissions for a smartwatch app?

Theme	Subtheme	Definition	Quote
Route	Direct through Settings, Apps and Permissions	Responses in this category changed permissions through following a direct route via the 'Settings' menu, then the 'Apps' menu (or the 'Apps and Notifications' menu) and then the 'Permissions' menu (or 'App Permissions' menu). These routes are equivalent, and only differ based on the operating system version. This path may be preceded by unlocking the watch if a screen lock is enabled. This path might also be extended by mentioning certain apps, their permissions and changing them.	<i>"Settings, again. Apps and Notifications, let's see. App Permissions"</i>
	Experienced technical issues	Responses in this category experienced technical difficulties in changing permissions. The participant followed the route of the 'Settings' menu, then the 'Apps' menu (or the 'Apps and Notifications' menu), but failed to reach the final menu.	<i>"There's System Apps. Hello? Have I just gone off it? It's just decided to take me back"</i>
	Indirect route	Responses in this category changed permissions, but via an indirect route. They succeeded in the task but did not go straight through the 'Settings' menu, the 'Apps' menu (or the 'Apps and Notifications' menu), and then the 'Permissions' menu (or 'App Permissions' menu).	<i>"Okay so I will pick one and so I'm clicking on an app but that's not what I want to do"</i>
Confidence	Certain	Responses in this category were reasonably certain over the process/route to change permissions. They did not display doubts over the path they were taking. Even if technical issues were encountered, they might be certain over the route they are attempting.	<i>"I'm hitting Settings, I'm going to Apps and Notifications, I'm going to App Permissions"</i>
	Uncertain	Responses in this category were less certain over the process/route to change permissions. They either doubted the path or guessed at certain stages.	<i>"Okay I think I would still be in Settings. And it should be Apps and Notifications"</i>

Could you please show me how to change the permissions for a smartwatch app? *Continued*

Had not changed app permissions before	Responses in this category admitted to having never changed their app permissions before.	<i>“So I haven’t done that”</i>
--	---	---------------------------------

Table A.49: How serious do you feel the action of your smartwatch data being accessed is?

Theme	Subtheme	Definition	Quote
Serious factors	Dislike the principle of data access	Responses in this category indicated that data access has seriousness, since they disliked the fact that parties intrude into their privacy.	<i>“They would intrude too much into my privacy”</i>
	Concerned about targeted advertising	Responses in this category indicated that data access has seriousness, since they are concerned about recommendations and advertising being targeted or tailored.	<i>“I guess from that you could target sales as well”</i>
	Concerned about social engineering	Responses in this category indicated that data access has seriousness, since they are concerned about social engineering techniques.	<i>“It then contributes to some bigger pool which is then used for social engineering techniques”</i>
	Apps can be unethical even if consented	Responses in this category indicated that data access has seriousness, since companies can be nefarious even if collection is consented.	<i>“When you do consent to these things...it’s not outside of the company’s remit to have nefarious uses”</i>
	Apps can leak data through incompetence	Responses in this category indicated that data access has seriousness, since apps can be vulnerable due to incompetent design.	<i>“Your app is vulnerable by itself. So it’s either nefarious or incompetent”</i>
	Dislike location being tracked	Responses in this category indicated that data access has seriousness, since they dislike their location being tracked.	<i>“If they could track where exactly I am at a moment then...I do not want to happen”</i>
	Privacy protection is not perfect	Responses in this category indicated that data access has seriousness, since they believe information is not necessarily protected.	<i>“I don’t think it actually does protect all of your data necessarily”</i>

How serious do you feel the action of your smartwatch data being accessed is? *Continued*

Contingent	Depends if consented	Responses in this category indicated that seriousness of data access is dependent on whether consent has been acquired. If the accessing party has received permission, the issue is deemed less serious.	<i>“Without consent of course I would be concerned”</i>
	Depends on data/sensitivity	Responses in this category indicated that seriousness is dependent on what kind of data is accessed and how sensitive it is.	<i>“It depends on what data is accessed”</i>
	Depends if app provides functionality/benefits	Responses in this category indicated that seriousness of data access is dependent on whether the app provides functionality in return. If access is expected to support a feature, it is more accepted. If data sharing leads to a benefit, it is also more accepted.	<i>“I need to check whether I get something from using my data or not”</i>
	Depends on how data is used	Responses in this category indicated that seriousness of data access is dependent on how the information is used. If the data is not used ethically, the issue is more serious.	<i>“I’m just for ethical use of data”</i>
	Depends on which company	Responses in this category indicated that seriousness of data access is dependent on which company undertakes this action. If the company is either trusted or expected to have access, then it is deemed less serious.	<i>“Depends on the company. There are some companies I trust more than others”</i>
Not serious	Accept/assume data will be accessed	Responses in this category indicated that data access lacks seriousness, since they assume technology uses information frequently.	<i>“So I’m working the assumption that the data from my watch and the Android software will collect it”</i>
	Access necessary to use apps	Responses in this category indicated that data access lacks seriousness, since apps will not work properly without data.	<i>“Even if you do enable these data restrictions...you’re just going to have apps that don’t work half the time”</i>
	Not concerned about advertising	Responses in this category indicated that data access lacks seriousness, since are not concerned about advertisements or recommendations.	<i>“I get targeted advertising for that...so I’m probably not too concerned about that”</i>

How serious do you feel the action of your smartwatch data being accessed is? *Continued*

	Do not have sensitive data	Responses in this category indicated that data access lacks seriousness, since the data on their watch is not sensitive.	<i>"I don't have many sensitive data"</i>
--	----------------------------	--	---

Table A.50: How effective do you think smartwatch settings can be in protecting data?

Theme	Subtheme	Definition	Quote
Effective in protection	Confident that settings protect data	Responses in this category are confident that settings can actually protect smartwatch data. This might include a brief description of how apps/watch loses access to certain pieces of data. Responses may simply be in praise of the effectiveness of the settings, even not completely perfect.	<i>"I'm quite confident I'd be able to make it a lot more secure"</i>
	Trust since can view source code	Responses in this category are confident in settings since they can check the source code themselves.	<i>"So the smartphone setting are really good, I mean...you can really check how Android is made as you can access the code"</i>
Partially effective	Reduces/limits amount of data access	Responses in this category believe settings are partially effective as they limit and protect the access to data.	<i>"It can help you limit the amount you're putting out there"</i>
	Helps feel more in control	Responses in this category believe settings are partially effective since they give users a greater sense of control.	<i>"It makes you feel a bit more in control"</i>
Contingent	Depends on which settings	Responses in this category believe the efficacy of settings protection depends on which settings are used. General watch settings might differ from app-specific options.	<i>"I mean there are different levels aren't there"</i>
	Depends on adversary sophistication	Responses in this category believe the efficacy of settings protection depends on how skilled the adversary is.	<i>"If somebody is really smart I think there are ways to do it"</i>
Not effective	Other routes into devices	Responses in this category doubt the efficacy of settings since there are always flaws and backdoors that allow the access of data.	<i>"Backdoors in all of our technological devices, so nothing is truly safe"</i>

How effective do you think smartwatch settings can be in protecting data? *Continued*

Can just steal the watch	Responses in this category doubt the efficacy of settings since people could just steal the smartwatch.	<i>“But people can still steal your watch”</i>
Could access through stolen phone	Responses in this category doubt the efficacy of settings since people could access data by stealing the paired smartphone.	<i>“There’s still a chance that people can steal both your phone and watch”</i>
Impractical as need to allow apps	Responses in this category doubt the efficacy of settings as apps require these settings to be loosened.	<i>“They’re not really effective because many apps don’t work if you don’t give them the authorisation”</i>
Do not provide meaningful control	Responses in this category doubt the efficacy of settings since they do not offer fine-grained control. For example, users cannot choose what data is used for or whether it is stored locally or sent to a server.	<i>“There’s a difference between...using only locally, or sending and we will send it to the server and store it”</i>
Requires some degree of knowledge	Responses in this category doubt the efficacy of settings since users must possess knowledge to recognise they are there.	<i>“I don’t think necessarily everybody who has a smartwatch would know that’s something they could do”</i>
Still agreed to Terms and Conditions	Responses in this category doubt the efficacy of settings since users still accept Terms and Conditions. Therefore, access to data has been consented even if the settings are adjusted.	<i>“You have to accept certain conditions and I guess that’s up to me”</i>
Privacy is not a salient thought	Responses in this category doubt the efficacy of settings since privacy is simple to overlook or forget.	<i>“I feel they can be very effective, but I feel it’s also easy to overlook them”</i>

Table A.51: How able do you feel *you* are to protect your smartwatch's data?

Theme	Subtheme	Subsubtheme	Definition	Quote
Confident	Confident in using settings		Responses in this category have confidence in protecting data, since they believe they can navigate through the settings. Answers concern self-efficacy, rather than the abilities of the settings.	<i>"I'm pretty confident that if I want to protect to certain parts of my data, I can"</i>
	Privacy interface is simple to use		Responses in this category have confidence in protecting data, since they believe the settings are easy to navigate.	<i>"Because it is quite easy interface to navigate"</i>
Sacrifices privacy for convenience			Responses in this category have confidence in protecting data, but often trade that data for functionality or convenience.	<i>"I do not do that because it would become useless at many things"</i>
Not confident	Not confident in settings	Settings do not protect everything	Responses in this category lack confidence in protecting data, since they are unsure whether settings protect everything.	<i>"My confidence in being able to operate these things is kinda predicated on the assume that they do exactly what they say"</i>
		Watches still monitor other things	Responses in this category lack confidence in protecting data, since the watch monitors information regardless of what you deactivate.	<i>"They're still looking at what's going on, even if you have told them that they can't"</i>
		Sophisticated attacks too hard	Responses in this category lack confidence in protecting data, since if a skilled adversary wants to access information, they will.	<i>"I think if someone really wanted to find out my data, I think they could"</i>
		Might have previously consented	Responses in this category lack confidence in protecting data, since users might have previously given permission. Therefore, even if a person changes their settings, companies may have already accessed that data.	<i>"I don't trust that I didn't just click Yes when I installed it"</i>

How able do you feel *you* are to protect your smartwatch's data? *Continued*

		Do not have complete control	Responses in this category lack confidence in protecting data, since participants do not feel they have complete control of their information.	<i>"I know I don't have absolute control over what happens"</i>
Not confident in self		Lacks experience with settings	Responses in this category lack confidence in protecting data, since they have not tried settings before.	<i>"Given that I haven't really tried that many steps to protect my privacy, I would say...not that confident"</i>
		Privacy lacks salience	Responses in this category lack confidence in protecting data, since privacy is easy to forget.	<i>"It's not the first thing I think of when I download it"</i>
		Would use apps regardless	Responses in this category lack confidence in protecting data, since the participant would continue to use an app even if there were no privacy options.	<i>"If I...thought, 'there's no option for whether or not they can share my data', I wouldn't then delete the app"</i>

Table A.52: Do you feel you receive benefits from using data-accessing apps?

Theme	Subtheme	Subsubtheme	Definition	Quote
Concrete advantages	Location access	Ordering taxis	Responses in this category receive the benefit of ordering Uber taxis by sharing their location with apps.	<i>"As I said, like Uber knows your location, sends you cars"</i>
		Run tracking	Responses in this category receive the benefit of tracking their running by sharing their location with apps.	<i>"So like the running app that I used, tracking your location"</i>
		Navigation through maps	Responses in this category receive the benefit of using maps for navigation by sharing their location with apps.	<i>"Google Maps, for example, receives my location"</i>
	Heart rate tracking from body sensor access	Responses in this category receive the benefit of heart rate tracking from granting sensor access to the Fit app.	<i>"I use the Fit app just to keep track of my heartrate"</i>	
	Messaging from contacts access	Responses in this category receive the benefit of messaging friends by sharing their contacts with apps.	<i>"The messaging service I downloaded, that would have had access to all my contacts"</i>	

Do you feel you receive benefits from using data-accessing apps? *Continued*

	Reminders from calendar access	Responses in this category receive the benefit of meeting reminders by sharing their calendar with apps.	<i>“Calendar notifications so I don’t miss any meetings”</i>
	Interesting altimeter readings	Responses in this category receive the benefit of altimeter recordings by sharing their barometer data with apps.	<i>“I have an altimeter on there which was more useful than helpful”</i>
Claimed advantages	Grant data to receive app benefits	Responses in this category claim that their main reason they share data is to gain benefits.	<i>“I think that’s the main reason we let them have our data, because they give us some sort of use”</i>
	Can see advantages of such apps	Responses in this category can imagine how benefits might be provided by data-accessing apps.	<i>“Maybe if you used Maps you could go to some place...then maybe it could recommend food places”</i>
Only allow access to data if the app is beneficial		Responses in this category indicate the user is willing to allow access to their data, but only if the apps provide benefits on a case-by-case basis.	<i>“It has to be decided on a case-by-case basis”</i>
Have not used many data-accessing apps		Responses in this category indicate that the participant is unsure of benefits as they haven’t used a wide range of apps.	<i>“I haven’t used all the different apps”</i>
Risks	App benefits are questionable	Responses in this category believe that benefits from data-accessing apps are either minor, unclear or questionable.	<i>“I guess the one you could say is a bit more questionable is the Google Fit”</i>
	Apps can be invasive	Responses in this category believe that data access by such apps can be invasive. This might relate to concerns over personal data.	<i>“I can see how it’s quite invasive sometimes”</i>

Table A.53: How much effort do you feel it is to protect your smartwatch's data?

Theme	Subtheme	Definition	Quote
Little effort	Settings are easy to use	Responses in this category support that it is little effort to protect data, since settings are quick and easy to use.	<i>"No effort at all, it's very easy"</i>
	Supported by watch manufacturer	Responses in this category support that it is little effort to protect data, since watch manufacturers (Huawei and Google) provide security through updates and settings.	<i>"Place some sort of trust in the system that the software updates do some sort of job"</i>
Contingent	Depends on user's knowledge	Responses in this category believe the effort required to protect data depends on how aware the person is of how to use settings.	<i>"I mean it just takes you knowing that it's a possibility"</i>
	Depends if you want app functionality	Responses in this category believe the effort required to protect data depends on how much functionality is desired. If no features are used, then settings need never be changed.	<i>"But if you want to sometimes use it...then you have to switch it on or off"</i>
	Depends if settings are private by default	Responses in this category believe the effort required to protect data depends on whether settings are protective by default.	<i>"It depend what the default settings are"</i>
Same difficulty as protecting smartphone data		Responses in this category believe the effort required to protect data is similar to that of protecting data on a smart phone.	<i>"I think it's probably at the same level as protecting data on a phone"</i>
Greater effort	More effort than most will invest	Responses in this category believe data protection requires greater effort, since the average user will wish to do the absolute minimum.	<i>"More than the average person would be happy to spend"</i>
	Requirement to protect watch hardware	Responses in this category believe data protection requires great effort, since users must take care of the physical watch.	<i>"Of course you need to take care of your watch"</i>
	Privacy is not a salient thought	Responses in this category believe data protection requires great effort, since privacy is not an immediate thought when you acquire a smartwatch.	<i>"You may not have that forethought to necessarily do that immediately"</i>
	Protecting privacy is boring	Responses in this category believe data protection requires great effort since managing settings is not an exciting task.	<i>"It's not like managing your own privacy settings is the most exciting thing"</i>

How much effort do you feel it is to protect your smartwatch's data? *Continued*

	Settings are complex	Responses in this category believe data protection requires great effort since the privacy settings are obscure or annoying to navigate.	<i>"The settings may be a little obscure"</i>
--	----------------------	--	---

Table A.54: What would lead you to use protective tools more often?

Theme	Subtheme	Definition	Quote
Different environment	High-crime environment	Responses in this category believe they would use protective tools more often if they entered a dangerous or higher-crime environment. This category concerns the risk of injury, not political suppression.	<i>"Thefts are common there, I would definitely want to have these settings"</i>
	Country with less freedom	Responses in this category believe they would use protective tools more often if they were in a repressive country with less individual freedom.	<i>"You're living in a repressed country, you don't want people figuring out who it is"</i>
App issues	App maker discovered to be bad	Responses in this category believe they would use protective tools more often if the maker of one of their apps was found to be making mistakes or acting maliciously.	<i>"One of the developers of the apps got into some sort of scandal"</i>
	App maker discovered to be hacked	Responses in this category believe they would use protective tools more often if the maker of one of their apps was hacked.	<i>"Some big app maker was hacked for personal data"</i>
	Acquired sensitive job	Responses in this category believe they would use protective tools more often if they acquired a job of a sensitive or secretive nature.	<i>"If say I was working in a sensitive area"</i>
	Stored more sensitive data on watch	Responses in this category believe they would use protective tools more often if they had more sensitive content on their watch.	<i>"If that was something that was, say more potentially sensitive...I would be more cognisant of using the features"</i>

What would lead you to use protective tools more often? *Continued*

More important that location was private	Responses in this category believe they would use protective tools more often if their location privacy was more important.	<i>"I'm a student in Oxford temporary, so even if they start tracking...where I go, it's literally no problem"</i>
Prompted by the watch interface	Responses in this category believe they would use protective tools more often if they were asked by the smartwatch.	<i>"If it asked me, I might have done it. But it didn't, so I didn't"</i>
Negative information released about watch	Responses in this category believe they would use protective tools more often if new information highlighted the risks of smartwatches.	<i>"If a new piece of information comes to the light...I would probably more careful"</i>
Suffered negative consequences	Responses in this category believe they would use protective tools more often if their personal data was accessed.	<i>"I mean, in the circumstance that someone did use my data"</i>
Leaving watch unattended	Responses in this category believe they would use protective tools more often if they were leaving their watch unattended while completing tasks.	<i>"Leaving your watch unattended potentially, definitely I would want to maximise my security"</i>

Table A.55: Why do you think the Privacy Paradox might occur?

Theme	Subtheme	Subsubtheme	Definition	Quote
Concern claims	Oppose principle but not practice		Responses in this category believe the Paradox occurs because people oppose the thought of data violation, rather than the act in practice.	<i>"It's the thought of people accessing your data without permission, rather than in practice"</i>
	Respond with certainty to questions		Responses in this category believe the Paradox occurs because people will respond to a question with a certain answer. Therefore, they will claim strong concern, even if their views are mild.	<i>"When you ask a person that specific question, they will instantly tell you that they will feel concerned"</i>

Why do you think the Privacy Paradox might occur? *Continued*

		Responses are an aspiration	Responses in this category believe the Paradox occurs because people aspire to ideally have perfect privacy, but with the least effort.	<i>“Ideally I would like to have as much privacy as I can with the least effort”</i>
Behaviour issues	Impediments	Protection requires knowledge	Responses in this category believe the Paradox occurs because protective behaviour requires a degree of knowledge.	<i>“I guess with knowledge comes power”</i>
		Effort to protect privacy	Responses in this category believe the Paradox occurs because protective behaviour requires annoying effort and individuals are often lazy or not bothered.	<i>“I just think people just can't be bothered”</i>
		Lax default settings	Responses in this category believe the Paradox occurs because data is often accessed through default settings.	<i>“A lot of the stuff would be by default”</i>
		Data required for services	Responses in this category believe the Paradox occurs because data access must be allowed for apps to work.	<i>“Sometimes apps don't work if you don't give them permissions”</i>
		Privacy vs convenience	Responses in this category believe the Paradox occurs because privacy is traded off against convenience, and often functionality is considered more desirable.	<i>“I think it's the balance between convenience and privacy”</i>
		Norms	Reality of modern world	Responses in this category believe the Paradox occurs because we accept that our data is accessed as part of modern technology.
		Social norm to share data	Responses in this category believe the Paradox occurs because it is considered fashionable to use data-accessing technologies.	<i>“And also, using technology is kind of a status symbol”</i>

Why do you think the Privacy Paradox might occur? *Continued*

	Privacy can make stand out	Responses in this category believe the Paradox occurs because efforts to protect privacy can make an individual stand out.	<i>“If you’re the only person...who uses it, then you’re all you’re doing is pointing a big finger at yourself”</i>
	Privacy lacks salience	Responses in this category believe the Paradox occurs because people do not think about privacy in daily life.	<i>“You don’t immediately think, “how can this be used to track my data?””</i>
Data	Data thought not sensitive	Responses in this category believe the Paradox occurs because people do not regard their data as sensitive/important, and hence do not make effort to protect it.	<i>“I don’t have sensitive data so I think “okay, that means I don’t care””</i>
	Not bothered about data	Responses in this category believe the Paradox occurs because, deep down, people do not care about some pieces of data being accessed.	<i>“You don’t really care”</i>
	Rather pay with data	Responses in this category believe the Paradox occurs because people pay for services with their data, and would oppose funding companies through a subscription.	<i>“If suddenly you had to choose between a monthly subscription...or having this information being your payment...I think most people would actually fold”</i>
	Assumption data is private	Responses in this category believe the Paradox occurs because people take for granted that the default is that their data is private.	<i>“I think people kinda assume that the default is going to be in your interest”</i>
	Assumption data is needed	Responses in this category believe the Paradox occurs because people assume that if an app requests data, it needs it for a responsible function.	<i>“If it’s asked for this permission so it probably needs it for something”</i>

Why do you think the Privacy Paradox might occur? *Continued*

	Consequences are intangible / not immediate	Responses in this category believe the Paradox occurs because people do not directly experience or see their data being accessed. The consequences are also not seen at the time, and might be unknown.	<i>“You read that’s it’s an issue, but you don’t experience directly”</i>
Responses	Would take action if data was sensitive	Responses in this category indicate that they would take action if the accessed data was more sensitive.	<i>“If I would have sensitive data, I would definitely put some safeguards in place”</i>
	Easy to change smartwatch settings	Responses in this category believe that smartwatch privacy settings are easy to use.	<i>“But I think for the smartwatch it should be pretty easy”</i>
	Game highlighted privacy issues	Responses in this category highlighted that the individual thought more about privacy after playing the game app.	<i>“I think probably after playing the smartwatch app, after that I actually did look at apps”</i>

Table A.56: Why do you indicate concerns but not use settings to protect your data?

Theme	Subtheme	Subsubtheme	Definition	Quote
Concern claims	Not concerned enough to act		Responses in this category claim they had the Paradox as they were not actually concerned enough to take action.	<i>“I suppose my casual concern about it was not quite acute enough”</i>
	Oppose access in principle not practice		Responses in this category claim they had the Paradox as they are concerned about the principle, but do little in practice.	<i>“I am concerned, but I do minimum things to say that I’m concerned”</i>
Behaviour issues	Impediments	Too much effort	Responses in this category claim they had the Paradox as they were too lazy to take action.	<i>“Sometimes it was laziness, I’m not gonna lie”</i>
		Was not prompted	Responses in this category claim they had the Paradox as they were not prompted with opportunities to protect their privacy.	<i>“Why didn’t I use the privacy settings? Again, as I alluded to, it didn’t ask me”</i>

Why do you indicate concerns but not use settings to protect your data? *Continued*

	Privacy vs convenience	Responses in this category claim they had the Paradox since they traded off privacy against convenience, and often found functionality was more desirable.	<i>"So I'm constantly doing this trade-off and usually apps win"</i>
Norms	Privacy lacks salience	Responses in this category claim they had the Paradox as privacy is unseen, simple to overlook and easy to forget.	<i>"I think it's very much that you don't think about it"</i>
	People are passive	Responses in this category claim they had the Paradox as they believe most people defer action until they face an active issue.	<i>"I think human beings as a rule are generally quite passive sometimes"</i>
	Privacy is low priority	Responses in this category claim they had the Paradox as privacy was a lower priority than other issues.	<i>"It kinda gets shunted to the back of the list of things to do"</i>
	Privacy risk not in media	Responses in this category claim they had the Paradox as they were no issues in the media to encourage them action was necessary.	<i>"So because they're not in the media at the moment"</i>
Perceptions	Did not perceive risk	Responses in this category claim they had the Paradox as they did not think their watch was compromised or they were at risk of damage.	<i>"I did not have the feeling that my information was at current risk"</i>
	Data not considered sensitive	Responses in this category claim they had the Paradox as they did not believe their data was sensitive, important or of value to another party.	<i>"I don't have any sensitive data and I don't think they can damage me"</i>
	Privacy is inconvenient	Responses in this category claim they had the Paradox as they perceived protecting privacy as being inconvenient.	<i>"Having to put in a password every time I wanted to access a feature"</i>
	Nothing to hide	Responses in this category claim they had the Paradox as they did not believe they had anything to hide.	<i>"I have nothing important to hide"</i>

Why do you indicate concerns but not use settings to protect your data? *Continued*

		Data accessible in other ways	Responses in this category claim they had the Paradox as they were not incentivised to act, since their data can be accessed in many other ways.	<i>"There are many other ways people can access a similar degree of personal data"</i>
		Data access limited to phone	Responses in this category claim they had the Paradox as they did not protect data since access was limited to that which was on the phone.	<i>"The worst it can get is access to everything on my phone"</i>
Alternative protection	Protecting physical access to watch		Responses in this category claim they guarded their data by focusing on physically protecting the watch.	<i>"I placed the emphasis on "okay, I'll look after the watch"</i>
	Limiting the apps installed		Responses in this category claim they protected privacy by limiting which apps they installed.	<i>"I could say that I intentionally limit the amount of apps I have on the smartwatch"</i>
Responses	Would act if under greater risk		Responses in this category claim they would check protective approaches if they thought they were under greater risk or their data was more sensitive.	<i>"If I think it can then I will dive into privacy settings"</i>
Other thoughts	Learned about privacy through game		Responses in this category believed they learned about privacy through the smartwatch game.	<i>"Before playing the game I didn't even though I knew there's a feature"</i>
	Easier to protect privacy on smartwatch		Responses in this category believed that smartwatch settings were easier to protect privacy.	<i>"The methods of trying to protect this are somewhat more accessible"</i>

Appendix B

Smartwatch Game Images

B.1 Online Prototype Game



Figure B.1: Name Selection Menu

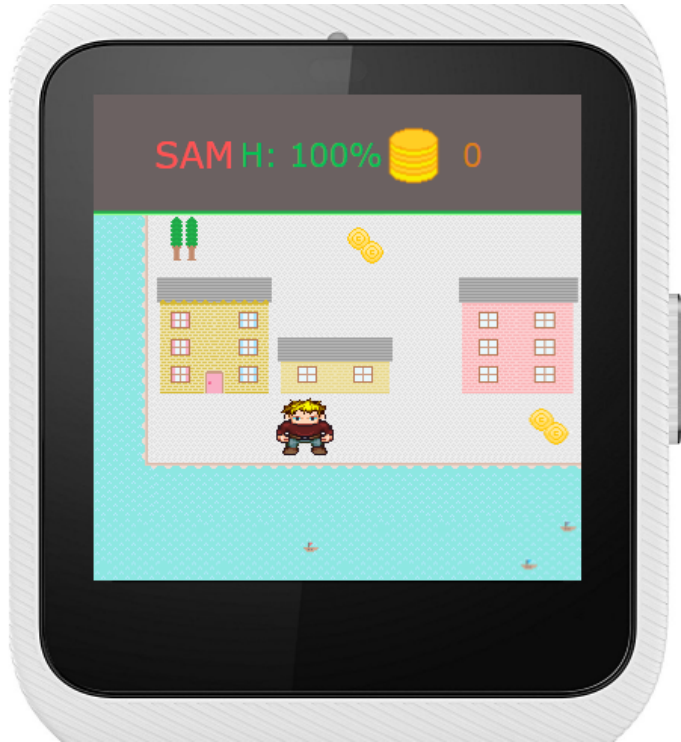


Figure B.2: Game Map



Figure B.3: Privacy Challenge: Disable Permissions



Figure B.4: Privacy Challenge: Simulated Settings




Figure B.5: Privacy Challenge: Task Completed



Figure B.6: Privacy Question: Enable Passcode

B.2 Wear OS Game



Figure B.7: Huawei Watch 2 Device



Figure B.8: Avatar Selection Menu



Figure B.9: Game Map: Morning (Default) Difficulty



Figure B.10: Game Map: Night (Highest) Difficulty



Figure B.11: Privacy Challenge: Disable GPS



Figure B.12: Privacy Challenge: Simulated Settings

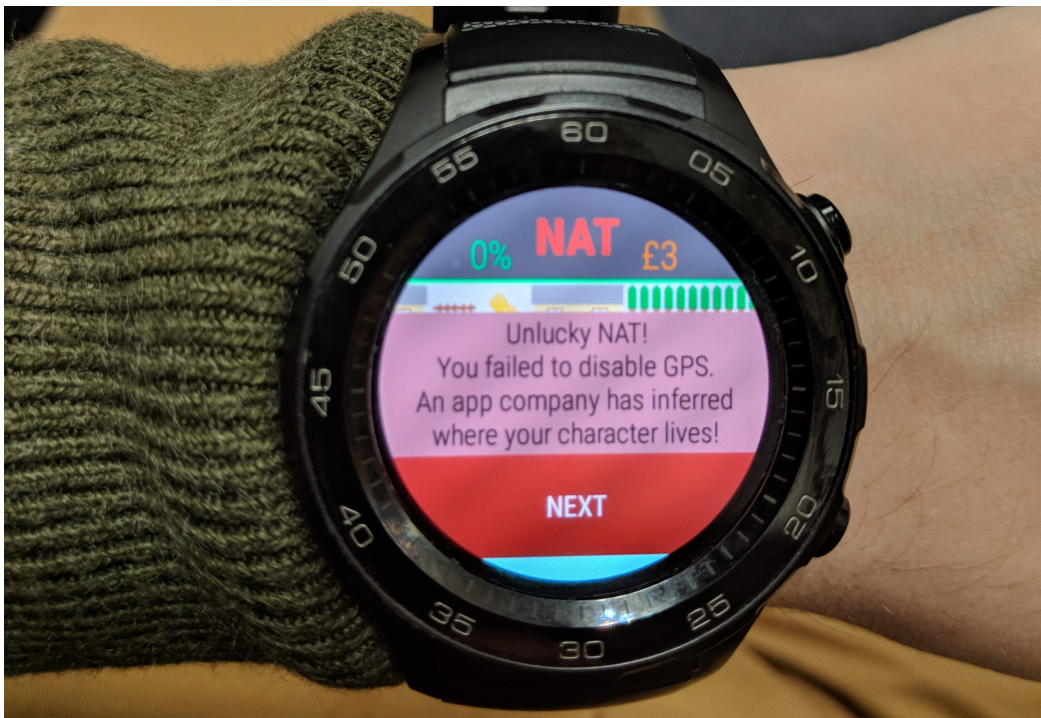


Figure B.13: Privacy Challenge: Task Failure



Figure B.14: Privacy Question: Revoke Permissions

Appendix C

Additional Participant Profiles

Participant B

B sought to balance privacy against functionality. Data access was deemed severe if messages were read; not so serious if it was browsing data. While they did not have an extensive knowledge of the risks, they recognised that Google used information. Functionality benefits were enjoyed, particularly run tracking and heart-rate monitoring. However, when privacy settings were used, they appreciated that their data was safe. They were also confident in using these features, but accepted that they impeded functionality. Protection effort was thought to be low; similar to that of guarding a smartphone. The participant justified the Paradox on data being weighed against convenience. Although they exchanged some details for functionality, they made informed and balanced decisions.

Participant E

E expressed how privacy is easily forgotten in everyday life. But when the issue was highlighted, they acted to protect their data. They deemed access to be severe when messages were targeted, but less serious for fitness data. The participant knew that apps could read smartwatch details, but trusted that most parties had good intentions. They received benefits from such access, recognising that WhatsApp would be useless without contacts. Privacy settings were partially trusted, but skilled adversaries were considered too powerful. The user had some confidence in using protection, and believed these tools were quick to configure. Although they claimed to often forget about privacy, they enabled a screen lock after playing the game. The Privacy Paradox was justified on protection being a low priority compared to other tasks.

Participant F

F was concerned about privacy, but admitted to not investigating protection. Data access was considered severe if unconsented, and still serious if used for nefarious purposes. However, they doubted they were highly vulnerable, since they installed few apps. For this reason, the participant received benefits from only fitness and messaging applications. After they learned about the privacy settings, they had confidence they could adjust these tools. Nevertheless, they did not believe they offered absolute protection. The user thought settings were little effort, but accepted that some knowledge was required. The Privacy Paradox was explained by many people lacking awareness of protection.

Participant H

H doubted that their watch's data was highly sensitive. However, when presented with potential issues, they expressed concern. Data access was not deemed severe, but potentially annoying if used for location tracking. Although they had some awareness of information collection, they did not know exact details. The user did not install many apps, and therefore was limited in the benefits they experienced. When discussing privacy protection, they lacked confidence. They believed settings were helpful but could be overcome by skilled attackers. They also doubted their own abilities, since they had not used the tools before. The participant expressed that their motivation was driven by data sensitivity. Therefore, even now they know of the settings, they would be unlikely to use them frequently. They justified the Privacy Paradox on people not being bothered about protection.

Participant J

J did not deem their data to be particularly sensitive. While access was severe if done by untrusted companies, they were less concerned if a service was provided. They also saw little vulnerability, since they were not a person of interest. The participant received navigational benefits from Google Maps, and considered this to be a fair transaction. They felt skilled in using privacy settings, and regarded them as little effort. However, they believed external hackers might overcome smartwatch protection. They justified the Privacy Paradox on people failing to consider the negative consequences. If the individual was aware of a risk, it is likely they would investigate their settings.

Bibliography

- [1] M S Ackerman and L Cranor. Privacy critics: UI components to safeguard users' privacy. In *Proceedings of CHI'99 Extended Abstracts on Human Factors in Computing Systems*, pages 258–259, 1999.
- [2] M S Ackerman, L F Cranor, and J Reagle. Privacy in e-commerce: Examining user scenarios and privacy preferences. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, pages 1–8, 1999.
- [3] A Acquisti, L Brandimarte, and G Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015.
- [4] A Acquisti and R Gross. Imagined communities: Awareness, information sharing, and privacy on the Facebook. *Privacy Enhancing Technologies in Lecture Notes in Computer Science*, 4258:36–58, 2006.
- [5] A Acquisti and J Grossklags. Losses, gains, and hyperbolic discounting: An experimental approach to information security attitudes and behavior. In *Proceedings of the 2nd Workshop on the Economics of Information Security*, pages 1–27, 2003.
- [6] A Acquisti and J Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1):26–33, 2005.
- [7] J G Adair. The Hawthorne effect: A reconsideration of the methodological artifact. *Journal of Applied Psychology*, 69(2):334–345, 1984.
- [8] I Adjerid, E Peer, and A Acquisti. Beyond the Privacy Paradox: Objective versus relative risk in privacy decision making. *MIS Quarterly*, 42(2):465–488, 2018.
- [9] M Adorjan and R Ricciardelli. A new privacy paradox? Youth agentic practices of privacy management despite ‘nothing to hide’ online. *Canadian Review of*

- Sociology*, 2018. [Currently online; volume/number awaiting physical publication].
- [10] E Aguirre, A L Roggeveen, D Grewal, and M Wetzels. The personalization-privacy paradox: Implications for new media. *Journal of Consumer Marketing*, 33(2), 2016.
- [11] O Ajao, J Hong, and W Liu. A survey of location inference techniques on Twitter. *Journal of Information Science*, 41(6):855–864, 2015.
- [12] I Ajzen. *Attitudes, personality, and behavior*. Dorsey Press, 1988.
- [13] I Ajzen. The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2):179–211, 1991.
- [14] I Ajzen and M Fishbein. Attitude-behavior relations: A theoretical analysis and review of empirical research. *Psychological Bulletin*, 84(5):888–918, 1977.
- [15] A Aktypi, J R C Nurse, and M Goldsmith. Unwinding Ariadne’s identity thread: Privacy risks with fitness trackers and online social networks. In *Proceedings of the 2017 Conference on Multimedia Privacy and Security*, pages 1–11, 2017.
- [16] M Al-Sharrah, A Salman, and I Ahmad. Watch your smartwatch. In *Proceedings of the 2018 International Conference on Computing Sciences and Engineering*, pages 1–5, 2018.
- [17] T Alashoor and R Baskerville. The privacy paradox: The role of cognitive absorption in the social networking activity. In *Proceedings of the 36th International Conference on Information Systems*, pages 1–20, 2015.
- [18] T Alashoor, M Keil, L Liu, and J Smith. How values shape concerns about privacy for self and others. In *Proceedings of the 36th International Conference on Information Systems*, 2015.
- [19] Y Albayram, M M H Khan, and M Fagan. A study on designing video tutorials for promoting security features: A case study in the context of Two-Factor Authentication (2FA). *International Journal of Human-Computer Interaction*, 33(11):927–942, 2017.

- [20] Y Albayram, M M H Khan, and T Jensen. "...better to use a lock screen than to worry about saving a few seconds of time": Effect of fear appeal in the context of smartphone locking behavior. In *Proceedings of the 13th Symposium on Usable Privacy and Security*, pages 49–63, 2017.
- [21] N Aleisa and K Renaud. Yes, I know this IoT device might invade my privacy, but I love it anyway! A study of Saudi Arabian perceptions. In *Proceedings of the 2nd International Conference on Internet of Things: Big Data and Security*, pages 198–205, 2017.
- [22] H Almuhiemedi, F Schaub, N Sadeh, I Adjerid, A Acquisti, J Gluck, L F Cranor, and Y Agarwal. Your location has been shared 5,398 times!: A field study on mobile app privacy nudging. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 787–796, 2015.
- [23] C Anderson. Presenting and evaluating qualitative research. *American Journal of Pharmaceutical Education*, 74(8), 2010.
- [24] D Andrews, B Nonnecke, and J Preece. Conducting research on the Internet: Online survey design, development and implementation guidelines. *International Journal of Human-Computer Interaction*, 16(2):185–210, 2003.
- [25] A Andrushevich, B Copigneaux, R Kistler, A Kurbatski, F Le Gall, and A Klaproth. Leveraging multi-domain links via the Internet of Things. *Internet of Things, Smart Spaces, and Next Generation Networking in Lecture Notes in Computer Science*, 8121:13–24, 2013.
- [26] L A Annetta. The "I's" have it: A framework for serious educational game design. *Review of General Psychology*, 14(2):105–112, 2010.
- [27] L A Annetta and S Holmes. Creating presence and community in a synchronous virtual learning environment using avatars. *International Journal of Instructional Technology and Distance Learning*, 3(8):27–43, 2006.
- [28] L A Annetta, M R Murray, S G Laird, and S C Bohr. Serious games: Incorporating video games in the classroom. *Educause Quarterly*, 29(3):16–22, 2006.
- [29] D A Armor and S E Taylor. When predictions fail: The dilemma of unrealistic optimism. In *Heuristics and Biases: The Psychology of Intuitive Judgment*, pages 334–347. Cambridge University Press, 2002.

- [30] A Aron and E N Aron. *Statistics for psychology*. Prentice-Hall, 1994.
- [31] I Arroyo, Y Liu, N Wixon, and S Schultz. Toward embodied game-based intelligent tutoring systems. *Intelligent Tutoring Systems in Lecture Notes in Computer Science*, 9684:488–490, 2016.
- [32] D Ashbrook and T Starner. Using GPS to learn significant locations and predict movement across multiple users. *Personal and Ubiquitous Computing*, 7(5):275–286, 2003.
- [33] A J Ayer. *Logical positivism*. Simon and Schuster, 1966.
- [34] T Baarslag, A T Alan, R Gomer, M Alam, C Perera, E H Gerding, and m c schraefel. An automated negotiation agent for permission management. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 380–390, 2017.
- [35] T Baarslag, A T Alan, R C Gomer, I Liccardi, H Marreiros, E H Gerding, and m c schraefel. Negotiation as an interaction mechanism for deciding app permissions. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pages 2012–2019, 2016.
- [36] T Baarslag, I Liccardi, E H Gerding, R Gomer, and m c schraefel. Negotiating mobile app permissions. In *Proceedings of the Amsterdam Privacy Conference*, 2015.
- [37] P Backlund and M Hendrix. Educational games - Are they worth the effort? A literature survey of the effectiveness of serious games. In *Proceedings of the 5th International Conference on Games and Virtual Worlds for Serious Applications*, 2013.
- [38] M Bada, A Sasse, and J R C Nurse. Cyber security awareness campaigns: Why do they fail to change behaviour? In *Proceedings of the International Conference on Cyber Security for Sustainable Society*, pages 118–131, 2015.
- [39] Y M Baek. Solving the privacy paradox: A counter-argument experimental approach. *Computers in Human Behavior*, 38:33–42, 2014.
- [40] I Baggili, J Oduro, K Anthony, F Breitingner, and G McGee. Watch what you wear: Preliminary forensic analysis of smart watches. In *Proceedings of the*

- 10th International Conference on Availability, Reliability and Security (ARES)*, pages 303–311, 2015.
- [41] M W Bailey. Seduction by technology: Why consumers opt out of privacy by buying into the Internet of Things. *Texas Law Review*, 94:1023, 2016.
- [42] M Balestrini, G Seiz, L L Peña, and G Camprodon. Onboarding communities to the IoT. *Internet Science in Lecture Notes in Computer Science*, 10673:19–27, 2017.
- [43] A Bandura. Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review*, 84(2):191–215, 1977.
- [44] L Barkhuus. The mismeasurement of privacy: Using contextual integrity to reconsider privacy in HCI. In *Proceedings of the 2012 SIGCHI Conference on Human Factors in Computing Systems*, pages 367–376, 2012.
- [45] S Barr. Are we all environmentalists now? Rhetoric and reality in environmental action. *Geoforum*, 35(2):231–249, 2004.
- [46] K L Barriball and A While. Collecting data using a semi-structured interview: A discussion paper. *Journal of Advanced Nursing*, 19(2):328–335, 1994.
- [47] S Barth and M de Jong. The privacy paradox - Investigating discrepancies between expressed privacy concerns and actual online behavior - A systematic literature review. *Telematics and Informatics*, 34(7):1038–1058, 2017.
- [48] L Baruh, E Secinti, and Z Cemalcilar. Online privacy concerns and privacy management: A meta-analytical review. *Journal of Communication*, 67(1):26–53, 2017.
- [49] M Bashir, C Hayes, A D Lambert, and J P Kesan. Online privacy and informed consent: The dilemma of information asymmetry. In *Proceedings of the 78th ASIS&T Annual Meeting: Information Science with Impact: Research in and for the Community*, 2015.
- [50] BBC News. Tor gets help to anonymise users of ‘dark web’. <https://www.bbc.co.uk/news/technology-34221869>, 2015. [Online; Accessed 15-Aug-2018].
- [51] K Becker. Commercial off-the-shelf games (COTS). In *Choosing and Using Digital Games in the Classroom*, pages 101–118. Springer, 2017.

- [52] R Benbunan-Fich. Usability of wearables without affordances. In *Proceedings of the 23rd Americas Conference on Information Systems*, pages 1–10, 2017.
- [53] P M Bentler and D G Bonett. Significance tests and goodness of fit in the analysis of covariance structures. *Psychological Bulletin*, 88(3):588–606, 1980.
- [54] A R Beresford, A Rice, N Skehin, and R Sohan. MockDroid: Trading privacy for application functionality on smartphones. In *Proceedings of the 12th Workshop on Mobile Computing Systems and Applications*, pages 49–54, 2011.
- [55] S Beuker. *Privacy paradox: Factors influencing disclosure of personal information among German and Dutch SNS users*. Master’s Thesis, University of Twente, 2016.
- [56] M Bloor. On the analysis of observational data: A discussion of the worth and uses of inductive techniques and respondent validation. *Sociology*, 12(3):545–552, 1978.
- [57] J Bonneau and S Preibusch. The privacy jungle: On the market for data protection in social networks. In *Economics of Information Security and Privacy*, pages 121–167. Springer, 2010.
- [58] H N Boone Jr. and D A Boone. Analyzing Likert data. *Journal of Extension*, 50(2), 2012.
- [59] C Bösch, B Erb, F Kargl, H Kopp, and S Pfattheicher. Tales from the dark side: Privacy dark strategies and privacy dark patterns. *Proceedings on Privacy Enhancing Technologies*, 2016(4):237–254, 2015.
- [60] A Botta, W De Donato, V Persico, and A Pescapé. Integration of cloud computing and Internet of Things: A survey. *Future Generation Computer Systems*, 56:684–700, 2016.
- [61] R E Boyatzis. *Transforming qualitative information: Thematic analysis and code development*. Sage, 1998.
- [62] L Brandimarte, A Acquisti, and G Loewenstein. Misplaced confidences privacy and the control paradox. *Social Psychological and Personality Science*, 4(3):340–347, 2010.

- [63] R Brandom. Samsung’s fridge of the future will let you check spoilage from your phone. The Verge, <https://www.theverge.com/2016/1/4/10707894/samsung-smart-refrigerator-connected-fridge-iot-ces-2016>, 2016. [Online; Accessed 15-Aug-2018].
- [64] V Braun and V Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101, 2006.
- [65] A Braunstein, L Granka, and J Staddon. Indirect content privacy surveys: Measuring privacy without asking about it. In *Proceedings of the 7th Symposium on Usable Privacy and Security*, 2011.
- [66] P Briggs, D Jeske, and L Coventry. Behavior change interventions for cybersecurity. In *Behavior Change Research and Theory*, pages 115–136. Academic Press, 2017.
- [67] J Brill. The Internet of Things: Building trust and maximizing benefits through consumer control. *Fordham Law Review*, 83(1):205–217, 2014.
- [68] B Brown. Studying the Internet experience. HP Laboratories Technical Report, <http://www.hp1.hp.com/techreports/2001/HPL-2001-49.pdf>, 2001. [Online; Accessed 15-Aug-2018].
- [69] T Brown and L Smith. *Reductionism and the development of knowledge*. Psychology Press, 2003.
- [70] S Bruyneel and S Dewitte. Health nudges: How behavioural engineering can reduce chocolate consumption. In *The Economics of Chocolate*, pages 157–170. Oxford University Press, 2016.
- [71] A Bryant. Re-grounding grounded theory. *Journal of Information Technology Theory and Application*, 4(1):25–42, 2002.
- [72] M Büchi, N Just, and M Latzer. Caring is not enough: The importance of Internet skills for online privacy protection. *Information, Communication & Society*, 20(8):1261–1278, 2017.
- [73] M Buhrmester and T Kwang. Amazon’s Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6(1):3–5, 2011.

- [74] J K Burgoon. Privacy and communication. *Annals of the International Communication Association*, 6(1):206–249, 1982.
- [75] C Cadwalladr and E Graham-Harrison. Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach. *The Guardian*, <https://www.forbes.com/sites/paullamkin/2018/02/22/smartwatch-popularity-booms-with-fitness-trackers-on-the-slide/>, 2018. [Online; Accessed 15-Aug-2018].
- [76] C F Camerer and G Loewenstein. Behavioural economics: Past, present and future. In *Advances in Behavioral Economics*, pages 3–51. Princeton University Press, 2011.
- [77] J Card. Anonymity is the Internet’s next big battleground. *The Guardian*, <https://www.theguardian.com/media-network/2015/jun/22/anonymity-internet-battleground-data-advertisers-marketers>, 2015. [Online; Accessed 15-Aug-2018].
- [78] J Carifio and R Perla. Resolving the 50-year debate around using and misusing Likert scales. *Medical Education*, 42(12):1150–1152, 2008.
- [79] M Carter, J Downs, B Nansen, and M Harrop. Paradigms of games research in HCI: A review of 10 years of research at CHI. In *Proceedings of the 1st ACM SIGCHI Annual Symposium on Computer-Human Interaction in Play*, pages 27–36, 2014.
- [80] E J Caruana, M Roman, J Hernández-Sánchez, and P Solli. Longitudinal studies. *Journal of Thoracic Disease*, 7(11):537–540, 2015.
- [81] J Casano, H Tee, J Agapito, I Arroyo, and M M T Rodrigo. Migration and evaluation of a framework for developing embodied cognition learning games. In *Proceedings of the 3rd Asia-Europe Symposium on Simulation & Serious Gaming*, pages 199–203, 2016.
- [82] A Cavoukian. Privacy by design: The 7 foundational principles. www.ipc.on.ca/wp-content/uploads/Resources/7foundationalprinciples.pdf, 2009. [Online; Accessed 15-Aug-2018].
- [83] A Cavoukian. Privacy by Design in law, policy and practice. <http://www.ontla.on.ca/library/repository/mon/25008/312239.pdf>, 2011. [Online; Accessed 15-Aug-2018].

- [84] R Cellan-Jones. Who is smart about TV - Samsung, Sony or Google? BBC News, <http://www.bbc.co.uk/news/technology-16483715>, 2012. [Online; Accessed 15-Aug-2018].
- [85] B S Chaparro, M H Phan, and J R Jardina. Usability and performance of tablet keyboards: Microsoft Surface vs. Apple iPad. In *Proceedings of the Human Factors and Ergonomics Society 57th Annual Meeting*, 2013.
- [86] T Chenoweth, R Minch, and T Gattiker. Application of Protection Motivation Theory to adoption of protective technologies. In *Proceedings of the 42nd Hawaii International Conference on System Sciences*, 2009.
- [87] H Cho, J Lee, and S Chung. Optimistic bias about online privacy risks: Testing the moderating effects of perceived controllability and prior experience. *Computers in Human Behavior*, 26(5):987–995, 2010.
- [88] B C K Choi and A W P Pak. A catalog of biases in questionnaires. *Preventing Chronic Disease*, 2(1), 2005.
- [89] E Christofides, A Muise, and S Desmarais. Hey mom, what’s on your Facebook? Comparing Facebook disclosure and privacy in adolescents and adults. *Social Psychological and Personality Science*, 3(1):48–54, 2012.
- [90] J Chun, A Dey, K Lee, and S Kim. A qualitative study of smartwatch usage and its usability. *Human Factors and Ergonomics in Manufacturing and Service Industries*, 28(4):186–199, 2018.
- [91] J Clark, P C van Oorschot, and C Adams. Usability of anonymous web browsing: An examination of tor interfaces and deployability. In *Proceedings of the 3rd Symposium on Usable Privacy and Security*, pages 41–51, 2007.
- [92] R Clarke. Introduction to dataveillance and information privacy, and definitions of terms. <http://www.rogerclarke.com/DV/Intro.html#Priv>, 1999. [Online; Accessed 15-Aug-2018].
- [93] J Cohen. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46, 1960.
- [94] J Cohen. *Statistical power analysis for the behavioral sciences*. Academic Press, 1977.

- [95] J Cole. How unreasonable searches of private documents caused the American Revolution. <http://www.juancole.com/2013/07/unreasonable-documents-revolution.html>, 2013. [Online; Accessed 15-Aug-2018].
- [96] D Collingridge. Validating a questionnaire. Sage, <https://www.methodspace.com/validating-a-questionnaire/>, 2014. [Online; Accessed 15-Aug-2018].
- [97] Collins English Dictionary. Smartwatch definition and meaning. <https://www.collinsdictionary.com/dictionary/english/smartwatch>, 2017. [Online; Accessed 15-Aug-2018].
- [98] L Columbus. Internet of Things market to reach \$267b by 2020. Forbes, <https://www.forbes.com/sites/louiscolumbus/2017/01/29/internet-of-things-market-to-reach-267b-by-2020/>, 2017. [Online; Accessed 15-Aug-2018].
- [99] Comres. Trust and confidence in data survey. Information Commissioner’s Office, http://www.comresglobal.com/wp-content/uploads/2017/11/IC0_trust-and-confidence-in-data_2017-1.pdf, 2017. [Online; Accessed 15-Aug-2018].
- [100] B D Cone, C E Irvine, M F Thompson, and T D Nguyen. A video game for cyber security training and awareness. *Computers & Security*, 26(1):63–72, 2007.
- [101] T M Connolly, E A Boyle, E MacArthur, and T Hainey. A systematic literature review of empirical evidence on computer games and serious games. *Computers & Education*, 59(2):661–686, 2012.
- [102] Consumer Reports. 5 steps to protect your smart phone from theft or loss. <https://www.consumerreports.org/cro/2014/04/5-steps-to-protect-your-smart-phone-against-theft-or-loss/index.htm>, 2014. [Online; Accessed 15-Aug-2018].
- [103] K P L Coopamootoo and T Gross. Why privacy is all but forgotten: An empirical study of privacy and sharing attitude. *Proceedings on Privacy Enhancing Technologies (PoPETs)*, 2017(4):39–60, 2017.
- [104] L Coventry, P Briggs, J Blythe, and M Tran. Using behavioural insights to improve the public’s use of cyber security best practices. Technical Report,

- UK Government Office for Science, <http://nrl.northumbria.ac.uk/23903/1/14-835-cyber-security-behavioural-insights.pdf>, 2014. [Online; Accessed 15-Aug-2018].
- [105] H Cramer. *Mathematical methods of statistics*. Princeton University Press, 1946.
- [106] S Creese, M Goldsmith, J R C Nurse, and E Phillips. A data-reachability model for elucidating privacy and security risks related to the use of online social networks. In *Proceedings of the 11th International Conference on Trust, Security and Privacy in Computing and Communications*, pages 1124–1131, 2012.
- [107] M J Culnan and P K Armstrong. Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization Science*, 10(1):104–115, 1999.
- [108] J DaPonte. Your television is watching you: How information stored by Internet-connected home devices could be used against us. *Index on Censorship*, 45(1):88–90, 2016.
- [109] S Dash. *The intruders: Unreasonable searches and seizures from King John to John Ashcroft*. Rutgers University Press, 2004.
- [110] B Debatin, J P Lovejoy, A K Horn, and B N Hughes. Facebook and online privacy: Attitudes, behaviors, and unintended consequences. *Journal of Computer-Mediated Communication*, 15(1):83–108, 2009.
- [111] S Derikx, M de Reuver, and M Kroesen. Can privacy concerns for insurance of connected cars be compensated? *Electronic Markets*, 26(1):73–81, 2016.
- [112] B Desarnauts. Wristly Insider’s report #45. Medium, <https://medium.com/wristly-thoughts/one-year-in-and-only-now-are-we-getting-to-know-apple-watch-owners-db60d565d041>, 2016. [Online; Accessed 15-Aug-2018].
- [113] S Deterding, R Khaled, L E Nacke, and D Dixon. Gamification: Toward a definition. In *Proceedings of the CHI 2011 Gamification Workshop*, 2011.

- [114] A Deuker. Addressing the privacy paradox by expanded privacy awareness - The example of context-aware services. *Privacy and Identity Management for Life in IFIP Advances in Information and Communication Technology*, 320:275–283, 2009.
- [115] A J DeWitt and J Kuljis. Aligning usability and security: A usability study of Polaris. In *Proceedings of the 2nd Symposium on Usable Privacy and Security*, pages 1–7, 2006.
- [116] T Dienlin and S Trepte. Is the privacy paradox a relic of the past? An in-depth analysis of privacy attitudes and privacy behaviors. *European Journal of Social Psychology*, 45(3):285–297, 2015.
- [117] T Dinev and P Hart. An extended privacy calculus model for e-commerce transactions. *Information Systems Research*, 17(1):61–80, 2006.
- [118] Q Do, B Martini, and K K R Choo. Is the data on your wearable device secure? An Android Wear smartwatch case study. *Software: Practice and Experience*, 47(3):391–403, 2017.
- [119] P Dolan, M Hallsworth, D Halpern, D King, and I Vlaev. MINDSPACE: Influencing behaviour for public policy. Technical report, Institute of Government, London School of Economics and Political Science, 2010.
- [120] M S Donovan, J D Bransford, and J W Pellegrino. *How people learn: Bridging research and practice*. National Academies Press, 1999.
- [121] A Downs. An economic theory of political action in a democracy. *Journal of Political Economy*, 65(2):135–150, 1957.
- [122] G D’Souza and J E Phelps. The privacy paradox: The case of secondary disclosure. *Review of Marketing Science*, 7(1), 2009.
- [123] J Duffy. 8 Internet things that are not IoT. Network World, <http://www.networkworld.com/article/2378581/internet-of-things/8-internet-things-that-are-not-iot.html>, 2014. [Online; Accessed 15-Aug-2018].
- [124] O J Dunn. Multiple comparisons among means. *Journal of the American Statistical Association*, 56(293):52–64, 1961.

- [125] L Eadicicco. A new wave of gadgets can collect your personal information like never before. Business Insider, <http://www.businessinsider.com/privacy-fitness-trackers-smartwatches-2014-10>, 2014. [Online; Accessed 15-Aug-2018].
- [126] M Elkhodr, S Shahrestani, and H Cheung. A review of mobile location privacy in the Internet of Things. In *Proceedings of the 10th International Conference on ICT and Knowledge Engineering*, pages 266–272, 2012.
- [127] N B Ellison, C Steinfield, and C Lampe. The benefits of Facebook “friends”: Social capital and college students’ use of online social network sites. *Journal of Computer-Mediated Communication*, 12(4):1143–1168, 2007.
- [128] EMC. Consumer perceptions on security - do they still care? Technical Report, <http://www.emc.com/collateral/brochure/consumer-perceptions-on-security.pdf>, 2014. [Online; Accessed 15-Aug-2018].
- [129] D Evans. The Internet of Things: How the next evolution of the Internet is changing everything. Cisco, https://www.cisco.com/c/dam/en_us/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf, 2011. [Online; Accessed 15-Aug-2018].
- [130] J R Evans and A Mathur. The value of online surveys. *Internet Research*, 15(2):195–219, 2005.
- [131] M Evans and K Morley. Home gadgets open to hackers. The Telegraph, <https://www.telegraph.co.uk/news/2017/07/24/internet-things-will-leave-home-gadgets-vulnerable-hacks-senior/>, 2017. [Online; Accessed 15-Aug-2018].
- [132] D Feldman. *Civil liberties and human rights in England and Wales*. Oxford University Press, 2002.
- [133] A P Felt, S Egelman, and D Wagner. “I’ve got 99 problems, but vibration ain’t one”: A survey of smartphone users’ concerns. In *Proceedings of the 2nd ACM Workshop on Security and Privacy in Smartphones and Mobile Devices*, pages 33–44, 2012.
- [134] N G Fielding. Self-report study. In *The SAGE Dictionary of Social Research Methods*, pages 276–277. Sage, 2006.

- [135] J Finch. The vignette technique in survey research. *Sociology*, 21(1):105–114, 1987.
- [136] M Fishbein. A theory of reasoned action: Some applications and implications. *Nebraska Symposium on Motivation*, 27:65–116, 1979.
- [137] R A Fisher. On the interpretation of X^2 from contingency tables, and the calculation of p. *Journal of the Royal Statistical Society*, 85(1):87–94, 1922.
- [138] R J Fisher. Social desirability bias and the validity of indirect questioning. *Journal of Consumer Research*, 20(2):303–315, 1993.
- [139] U Flick. Triangulation in qualitative research. In *A Companion to Qualitative Research*, pages 178–183. Sage, 2004.
- [140] D L Floyd, S Prentice-Dunn, and R W Rogers. A meta-analysis of research on protection motivation theory. *Journal of Applied Social Psychology*, 30(2):407–429, 2000.
- [141] B J Fogg. *Persuasive technology: Using computers to change what we think and do*. Morgan Kaufmann, 2003.
- [142] M Ford and W Palmer. Alexa, are you listening to me? An analysis of Alexa voice service network traffic. *Personal and Ubiquitous Computing*, 2018. [Currently online; volume/number awaiting physical publication].
- [143] S Foster, I Walker, R Feeney-Barry, and E Boyd. Barriers and benefits of home energy controller integration. UK Department for Business, Energy and Industrial Strategy, https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/545262/DECC_Barriers_and_Benefits_of_Home_Energy_Controllers_-_Final_report__1_.pdf, 2016. [Online; Accessed 15-Aug-2018].
- [144] C Franzwa, Y Tang, and A Johnson. Serious game design: Motivating students through a balance of fun and learning. In *Proceedings of the 5th International Conference on Games and Virtual Worlds for Serious Applications*, pages 1–7, 2013.
- [145] M Friedman. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American Statistical Association*, 32(200):675–701, 1937.

- [146] P Friess. *Digitising the industry: Internet of Things connecting the physical, digital and virtual worlds*. River Publishers, 2016.
- [147] N K Gale, G Heath, E Cameron, S Rashid, and S Redwood. Using the framework method for the analysis of qualitative data in multi-disciplinary health research. *BMC Medical Research Methodology*, 13(1), 2013.
- [148] A Gambino, J Kim, S S Sundar, J Ge, and M B Rosson. User disbelief in privacy paradox: Heuristics that determine disclosure. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pages 2837–2843, 2016.
- [149] L Garcia and F Quek. Qualitative research in information systems: Time to be subjective? In *Information Systems and Qualitative Research*, pages 444–465. Springer, 1997.
- [150] V Garg and J Camp. End user perception of online risk under uncertainty. In *Proceedings of the 45th Hawaii International Conference on System Science*, pages 3278–3287, 2012.
- [151] M Gašparović, P Nicolau, A Marques, C Silva, and L Marcelino. On privacy in user tracking mobile applications. In *Proceedings of the 11th Iberian Conference on Information Systems and Technologies*, pages 1–6, 2016.
- [152] J P Gee. *Situated language and learning: A critique of traditional schooling*. Psychology Press, 2004.
- [153] N Gerber, P Gerber, and M Volkamer. Explaining the Privacy Paradox - A systematic review of literature investigating privacy attitude and behavior. *Computers & Security*, 77:226–261, 2018.
- [154] S Gibbs. Samsung smart TVs send unencrypted voice recognition data across Internet. The Guardian, <https://www.theguardian.com/technology/2015/feb/19/samsung-smart-tvs-send-unencrypted-voice-recognition-data-across-internet>, 2015. [Online; Accessed 15-Aug-2018].
- [155] B Glanville. 72% of Brits haven't heard about GDPR. YouGov Omnibus, <https://yougov.co.uk/news/2018/03/01/72-brits-havent-heard-about-gdpr/>, 2018. [Online; Accessed 15-Aug-2018].

- [156] J Glenday. #DeleteFacebook movement builds up steam. The Drum, <https://www.thedrum.com/news/2018/06/18/deletefacebook-movement-builds-up-steam>, 2018. [Online; Accessed 15-Aug-2018].
- [157] J L Gómez-Barroso, C Feijóo, and I J Martínez-Martínez. Privacy calculus: Factors that influence the perception of benefit. *El Profesional de la Información*, 27(2):341–348, 2018.
- [158] N S Good and A Krekelberg. Usability and privacy: A study of Kazaa P2P file-sharing. In *Proceedings of the 2003 SIGCHI Conference on Human Factors in Computing Systems*, pages 137–144, 2003.
- [159] L D Goodwin and N L Leech. Understanding correlation: Factors that affect the size of r. *The Journal of Experimental Education*, 74(3):249–266, 2006.
- [160] R W Grant and J Sugarman. Ethics in human subjects research: Do incentives matter? *Journal of Medicine and Philosophy*, 26(9):717–738, 2004.
- [161] A Gruzd and Á Hernández-García. Privacy concerns and self-disclosure in private and public uses of social media. *Cyberpsychology, Behavior, and Social Networking*, 21(7):418–428, 2018.
- [162] C Hallam and G Zanella. Wearable device data and privacy: A study of perception and behavior. *World Journal of Management*, 7(1):82–91, 2016.
- [163] C Hallam and G Zanella. Online self-disclosure: The privacy paradox explained as a temporally discounted balance between concerns and rewards. *Computers in Human Behavior*, 68:217–227, 2017.
- [164] M H Hansen and W N Hurwitz. The problem of non-response in sample surveys. *Journal of the American Statistical Association*, 41(236):517–529, 1946.
- [165] M Harbach, M Hettig, S Weber, and M Smith. Using personal examples to improve risk communication for security & privacy decisions. In *Proceedings of the 32nd Annual ACM Conference on Human factors in Computing Systems*, pages 2647–2656, 2014.
- [166] E Hargittai and A Marwick. “What can I really do?”: Explaining the Privacy Paradox with online apathy. *International Journal of Communication*, 10:3737–3757, 2016.

- [167] Harris Insights & Analytics. IBM cybersecurity and privacy research. <http://newsroom.ibm.com/download/IBM+Cybersecurity+PR+Research+-+Final.pdf>, 2018. [Online; Accessed 15-Aug-2018].
- [168] S A Haslam, S A H C McGarty, and C McGarty. *Doing psychology: An introduction to research methodology and statistics*. Sage, 1997.
- [169] M Helft. Amid mounting questions, Dropbox tops 500 million users and says growth continues apace. Forbes, <https://www.forbes.com/sites/miguelhelft/2016/03/07/amid-growing-questions-dropbox-tops-500-million-users-and-says-growth-continues-pace/>, 2016. [Online; Accessed 15-Aug-2018].
- [170] M Helweg-Larsen and J A Shepperd. Do moderators of the optimistic bias affect personal or target risk estimates? A review of the literature. *Personality and Social Psychology Review*, 5(1):74–95, 2001.
- [171] A Hern. Windows 10: Microsoft under attack over privacy. The Guardian, <https://www.theguardian.com/technology/2015/jul/31/windows-10-microsoft-faces-criticism-over-privacy-default-settings>, 2015. [Online; Accessed 15-Aug-2018].
- [172] A Hern. WannaCry, Petya, NotPetya: How ransomware hit the big time in 2017. The Guardian, <https://www.theguardian.com/technology/2017/dec/30/wannacry-petya-notpetya-ransomware>, 2017. [Online; Accessed 15-Aug-2018].
- [173] A Hern. Cybercrime: £130bn stolen from consumers in 2017, report says. The Guardian, <https://www.theguardian.com/technology/2018/jan/23/cybercrime-130bn-stolen-consumers-2017-report-victims-phishing-ransomware-online-hacking>, 2018. [Online; Accessed 15-Aug-2018].
- [174] A Hern. Fitness tracking app Strava gives away location of secret US army bases. The Guardian, <https://www.theguardian.com/world/2018/jan/28/fitness-tracking-app-gives-away-location-of-secret-us-army-bases>, 2018. [Online; Accessed 15-Aug-2018].
- [175] L Heshusius and K Ballard. *From positivism to interpretivism and beyond: Tales of transformation in educational and social research (the mind-body connection)*. Teachers College Press, 1996.

- [176] Hewlett Packard Enterprise. Internet of Things security study: Smartwatches. Technical Report, http://go.saas.hpe.com/1/28912/2015-07-20/3251bm/28912/69038/IoT_Research_Series_Smartwatches.pdf, 2014. [Online; Accessed 15-Aug-2018].
- [177] C Hochleitner, C Graf, D Unger, and M Tscheligi. Making devices trustworthy: Security and trust feedback in the Internet of Things. In *Proceedings of the 4th International Workshop on Security and Privacy in Spontaneous Interaction and Mobile Phone Use*, 2012.
- [178] J H Hoepman. Privacy design strategies. *IFIP Advances in Information and Communication Technology*, 428:446–459, 2014.
- [179] C P Hoffmann, C Lutz, and G Ranzini. Privacy cynicism: A new approach to the privacy paradox. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*, 10(4), 2016.
- [180] B Hoh, M Gruteser, H Xiong, and A Alrabady. Enhancing security and privacy in traffic-monitoring systems. *IEEE Pervasive Computing*, 5(4):38–46, 2006.
- [181] M Hölbl, T Welzer, and L Nemeč Zlatolas. Security and privacy related issues in the Internet of Things. *Frontiers in Artificial Intelligence and Applications*, 280:321–326, 2016.
- [182] J Holvast. History of privacy. In *The Future of Identity in the Information Society*, pages 13–42. Springer Berlin Heidelberg, 2009.
- [183] A M Horcher. A tap on the wrist: Security usability for wearables. In *Proceedings of the 2nd Workshop on Inclusive Privacy and Security*, 2015.
- [184] M B Hoy. The “Internet of Things”: What it is and what it means for libraries. *Medical Reference Services Quarterly*, 34(3):353–358, 2015.
- [185] H Y Huang and M Bashir. Is privacy a human right? An empirical examination in a global context. In *Proceedings of the 13th Annual Conference on Privacy, Security and Trust (PST)*, pages 77–84, 2015.
- [186] H Y Huang and M Bashir. Privacy by region: Evaluation online users’ privacy perceptions by geographical region. In *Proceedings of the 1st Future Technologies Conference*, 2016.

- [187] G Hughes, S Silver, and A Lewins. Organising survey open-ended question data for CAQDAS. QUIC Working Paper, <https://www.surrey.ac.uk/sites/default/files/OEQ-document-per-case-vs-document-per-question-June2010.pdf>, 2010. [Online; Accessed 15-Aug-2018].
- [188] T Hughes-Roberts and S Furnell. Privacy as a secondary goal problem: An experiment examining control. *Information & Computer Security*, 23(4):382 – 393, 2015.
- [189] T Hughes-Roberts and E Kani-Zabihi. On-line privacy behavior: Using user interfaces for salient factors. *Journal of Computer and Communications*, 2(4):220, 2014.
- [190] P Humby. Overview of the UK population: February 2016. UK Office for National Statistics, <https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/articles/overviewoftheukpopulation/february2016>, 2016. [Online; Accessed 15-Aug-2018].
- [191] IDC. Worldwide quarterly wearable device tracker. <https://www.idc.com/getdoc.jsp?containerId=prUS43642518>, 2018. [Online; Accessed 15-Aug-2018].
- [192] J Ilieva, S Baron, and N M Healey. Online surveys in marketing research: Pros and cons. *International Journal of Market Research*, 44(3):361, 2002.
- [193] International Telecommunication Union. ITU Internet Reports 2005: The Internet of Things. <https://www.itu.int/osg/spu/publications/internetofthings/>, 2005. [Online; Accessed 15-Aug-2018].
- [194] C B Jackson and Y Wang. Addressing the Privacy Paradox through personalized privacy notifications. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(2), 2018.
- [195] J Jackson, N Allum, and G Gaskell. Perceptions of risk in cyberspace. In *Trust and Crime in Information Societies*, pages 245–281. Edward Elgar, 2005.
- [196] H Jeong, H Kim, R Kim, U Lee, and Y Jeong. Smartwatch wearing behavior analysis: A longitudinal study. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(3):60, 2017.

- [197] M Johnson, S Egelman, and S M Bellovin. Facebook and privacy: It's complicated. In *Proceedings of the 8th Symposium on Usable Privacy and Security*, 2012.
- [198] R B Johnson and A J Onwuegbuzie. Mixed methods research: A research paradigm whose time has come. *Educational Researcher*, 33(7):14–26, 2004.
- [199] C Jordan, M Knapp, D Mitchell, and M Claypool. CounterMeasures: A game for teaching computer security. In *Proceedings of the 10th Annual Workshop on Network and Systems Support for Games*, 2011.
- [200] L S Kao and C E Green. Analysis of variance: Is there a difference in means and what does it mean? *Journal of Surgical Research*, 144(1):158–170, 2008.
- [201] W D Kearney and H A Kruger. Theorising on risk homeostasis in the context of information security behaviour. *Information & Computer Security*, 24(5):496–513, 2016.
- [202] F Kehr, T Kowatsch, D Wentzel, and E Fleisch. Blissfully ignorant: The effects of general privacy concerns, general institutional trust, and affect in the Privacy Calculus. *Information Systems Journal*, 25(6):607–635, 2015.
- [203] M J Keith, J S Babb, and P B Lowry. A longitudinal study of information privacy on mobile devices. In *Proceedings of the 47th Hawaii International Conference on System Sciences*, pages 3149–3158, 2014.
- [204] K Kelley and K J Preacher. On effect size. *Psychological Methods*, 17(2):137–152, 2012.
- [205] P G Kelley, J Bresee, L F Cranor, and R W Reeder. A nutrition label for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, 2009.
- [206] P G Kelley, L F Cranor, and N Sadeh. Privacy as part of the app decision-making process. In *Proceedings of the 2013 SIGCHI Conference on Human Factors in Computing Systems*, pages 3393–3402, 2013.
- [207] T Kelley and B Bertenthal. Tracking risky behavior on the web: Distinguishing between what users 'say' and 'do'. In *Proceedings of the 9th International Symposium on Human Aspects of Information Security and Assurance*, pages 204–214, 2015.

- [208] B Kenney. Smartwatches vs fitness bands: There is a difference. <https://smartwatches.org/learn/smartwatches-vs-fitness-bands-difference/>, 2014. [Online; Accessed 15-Aug-2018].
- [209] K Kiili. Digital game-based learning: Towards an experiential gaming model. *The Internet and Higher Education*, 8(1):13–24, 2005.
- [210] S Kim, S Lee, and J Han. StretchArms: Promoting stretching exercise with a smartwatch. *International Journal of Human-Computer Interaction*, 34(3):218–225, 2018.
- [211] G Kirchgässner. *Homo oeconomicus: The economic model of behaviour and its applications in economics and other social sciences*. Springer Science & Business Media, 2008.
- [212] R E Kirk. Experimental design. In *The Blackwell Encyclopedia of Sociology*, pages 1533–1537. Wiley, 2007.
- [213] E Koblentz. Kaspersky extends Android security app to smartwatches. Tech Republic, <https://www.techrepublic.com/article/kaspersky-extends-android-security-app-to-smartwatches/>, 2016. [Online; Accessed 15-Aug-2018].
- [214] S Kokolakis. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security*, 64:122–134, 2017.
- [215] T L Koreshoff, T Robertson, and T W Leong. Internet of Things: A review of literature and products. In *Proceedings of the 25th Australian Computer-Human Interaction Conference: Augmentation, Application, Innovation, Collaboration*, pages 335–344, 2013.
- [216] G Kortuem, F Kawsar, V Sundramoorthy, and D Fitton. Smart objects as building blocks for the Internet of Things. *IEEE Internet Computing*, 14(1):44–51, 2010.
- [217] M Kusters and J Van der Heijden. From mechanism to virtue: Evaluating Nudge theory. *Evaluation*, 21(3):276–291, 2015.

- [218] T Kowatsch and W Maass. Critical privacy factors of Internet of Things services: An empirical investigation with domain experts. *Knowledge and Technologies in Innovative Information Systems in Lecture Notes in Business Information Processing*, 129:200–211, 2012.
- [219] M J Kraemer and I Fléchaïs. Researching privacy in smart homes: A roadmap of future directions and research methods. In *Living in the Internet of Things: Cybersecurity of the IoT*, 2018.
- [220] H Krasnova, E Kolesnikova, and O Guenther. “It won’t happen to me!”: Self-disclosure in online social networks. In *Proceedings of the 15th Americas Conference on Information Systems*, pages 1–9, 2009.
- [221] K Krishnan. Internet of Things data access and the fear of the unknown. IBM Blogs, <http://www.ibmbigdatahub.com/blog/internet-things-data-access-and-fear-unknown>, 2016. [Online; Accessed 15-Aug-2018].
- [222] W H Kruskal and W A Wallis. Use of ranks in one-criterion variance analysis. *Journal of the American Statistical Association*, 47(260):583–621, 1952.
- [223] P Kumaraguru, J Cranshaw, A Acquisti, L F Cranor, J Hong, M A Blair, and T Pham. School of phish: A real-world evaluation of anti-phishing training. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, 2009.
- [224] N R Lackey and A L Wingate. In *Advanced Design in Nursing Research*, pages 375–387. Sage.
- [225] D Laibson. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478, 1997.
- [226] P Lamkin. Smartwatch popularity booms with fitness trackers on the slide. Forbes, <https://www.forbes.com/sites/paullamkin/2018/02/22/smartwatch-popularity-booms-with-fitness-trackers-on-the-slide/>, 2018. [Online; Accessed 15-Aug-2018].
- [227] M Langheinrich. Privacy by Design - Principles of privacy-aware ubiquitous systems. *International Conference on Ubiquitous Computing in Lecture Notes in Computer Science*, 2201:273–291, 2001.
- [228] N B Lasson. *The history and development of the Fourth Amendment to the United States Constitution*. De Capo Press, 1970.

- [229] R S Laufer and M Wolfe. Privacy as a concept and a social issue: A multidimensional developmental theory. *Journal of Social Issues*, 33(3):22–42, 1977.
- [230] S Lederer, J I Hong, A K Dey, and J A Landay. Personal privacy through understanding and action: Five pitfalls for designers. *Personal and Ubiquitous Computing*, 8(6):440–454, 2004.
- [231] I Lee and K Lee. The Internet of Things (IoT): Applications, investments, and challenges for enterprises. *Business Horizons*, 58(4):431–440, 2015.
- [232] L Lee, J H Lee, S Egelman, and D Wagner. Information disclosure concerns in the age of wearable computing. In *Proceedings of the 2016 Workshop on Usable Security*, 2016.
- [233] S Lee, H R Ha, J H Oh, and N Park. The impact of perceived privacy benefit and risk on consumers’ desire to use Internet of Things technology. In *Proceedings of the International Conference on Human Interface and the Management of Information*, pages 609–619, 2018.
- [234] L Lenzi. Mass-market adoption of the smart home: Are we there yet? <https://www.appliancedesign.com/articles/95439-mass-market-adoption-of-the-smart-home-are-we-there-yet>, 2017. [Online; Accessed 15-Aug-2018].
- [235] M Lesk. License creep. *IEEE Security & Privacy*, 13(6):85–88, 2015.
- [236] J Leyden. UK’s Get Safe Online? ‘No one cares’ - run the blockbuster ads instead. The Register, https://www.theregister.co.uk/2013/09/25/gets_safe_online_has_failed_to_change_behaviours/, 2013. [Online; Accessed 15-Aug-2018].
- [237] H Li, J Wu, Y Gao, and Y Shi. Examining individuals’ adoption of health-care wearable devices: An empirical study from Privacy Calculus perspective. *International Journal of Medical Informatics*, 88:8–17, 2016.
- [238] R Lindblad and T Sasivani. *Consumer perceptions on the privacy-invasiveness of in-feed advertisements*. Bachelor’s Thesis, Jonkoping University, 2017.
- [239] R Lipman. Online privacy and the invisible market for our data. *Penn State Law Review*, 120:777–806, 2016.

- [240] A E Liska. A critical examination of the causal structure of the Fishbein/Ajzen attitude-behavior model. *Social Psychology Quarterly*, 47(1):61–74, 1984.
- [241] F Liu, R Ramanath, N Sadeh, and N A Smith. A step towards usable privacy policy: Automatic alignment of privacy statements. In *Proceedings of the 25th International Conference on Computational Linguistics*, pages 884–894, 2014.
- [242] Y Liu, K P Gummadi, B Krishnamurthy, and A Mislove. Analyzing Facebook privacy settings: User expectations vs. reality. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement*, pages 61–70, 2011.
- [243] C M Lyles, L S Kay, N Crepaz, J H Herbst, W F Passin, A S Kim, S M Rama, S Thadiparthi, J B DeLuca, and M M Mullins. Best-evidence interventions: Findings from a systematic review of HIV behavioral interventions for US populations at high risk, 2000-2004. *American Journal of Public Health*, 97(1):133–143, 2007.
- [244] W E Mackay. Triggers and barriers to customizing software. In *Proceedings of the 1999 SIGCHI Conference on Human Factors in Computing Systems*, pages 153–160, 1991.
- [245] A Mackey and S M Gass. *Second language research: Methodology and design*. Lawrence Erlbaum Associates, 2005.
- [246] L M MacLean, M Meyer, and A Estable. Improving accuracy of transcripts in qualitative research. *Qualitative Health Research*, 14(1):113–123, 2004.
- [247] M Madejski, M Johnson, and S M Bellovin. A study of privacy settings errors in an online social network. In *Proceedings of the 2012 IEEE International Conference on Pervasive Computing and Communications Workshops*, pages 340–345, 2012.
- [248] H Maldonado, J R L Lee, S Brave, C Nass, H Nakajima, R Yamada, K Iwamura, and Y Morishima. We learn better together: Enhancing eLearning with emotional characters. In *Proceedings of the 2005 Conference on Computer Support for Collaborative Learning*, pages 408–417, 2005.
- [249] Z Mani and I Chouk. Drivers of consumers’ resistance to smart products. *Journal of Marketing Management*, 33(1):76–97, 2017.

- [250] H B Mann and D R Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics*, 18(1):50–60, 1947.
- [251] J Manyika, M Chui, P Bisson, J Woetzel, R Dobbs, J Bughin, and D Aharon. The Internet of Things: Mapping the value beyond the hype. McKinsey Global Institute, <https://goo.gl/ARXWNN>, 2015. [Online; Accessed 15-Aug-2018].
- [252] C Maple. Security and privacy in the Internet of Things. *Journal of Cyber Policy*, 2(2):155–184, 2017.
- [253] P Martin. That government Cyber Aware website has cost £6.37 per visit since it launched. <http://www.alphr.com/politics/1005065/that-government-cyber-aware-website-has-cost-637-per-visit-since-it-launched>, 2017. [Online; Accessed 15-Aug-2018].
- [254] A Marwick and E Hargittai. Nothing to hide, nothing to lose? Incentives and disincentives to sharing information with institutions online. *Information, Communication and Society*, 2018. [Currently online; volume/number awaiting physical publication].
- [255] A E Marwick and D Boyd. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society*, 13(1):114–133, 2011.
- [256] S Matthews. Smartwatch dangers - Are you the target? War on Identity Theft, <http://www.waronidtheft.org/smartwatch-dangers-are-you-the-target/>, 2016. [Online; Accessed 15-Aug-2018].
- [257] R E Mayer, A Mathias, and K Wetzell. Fostering understanding of multimedia messages through pre-training: Evidence for a two-stage theory of mental model construction. *Journal of Experimental Psychology: Applied*, 8(3):147–154, 2002.
- [258] R McCloud, M Lerski, J Park, and T T Brooks. Wearable IoT computing: Interface, emotions, wearer’s culture, and security/privacy concerns. In *Cyber-Assurance for the Internet of Things*, pages 177–186. Wiley, 2017.
- [259] J H McDonald. Multiple tests. In *Handbook of Biological Statistics*, pages 256–265. Sparky House Publishing, 2009.

- [260] Q McNemar. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157, 1947.
- [261] C M Medaglia and A Serbanati. An overview of privacy and security issues in the Internet of Things. In *The Internet of Things*, pages 389–395. Springer, 2010.
- [262] J Melthis, S Tang, P Yang, M Hanneghan, and C Carter. Topologies for combining the Internet of Things and serious games. *Journal of Intelligent & Fuzzy Systems*, 31(5):2685–2696, 2016.
- [263] M Menfors and F Fernstedt. *Geotagging in social media - Exploring the privacy paradox*. Bachelor’s Thesis, University of Borås, 2015.
- [264] S Mennicken, J Vermeulen, and E M Huang. From today’s augmented houses to tomorrow’s smart homes: New directions for home automation research. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 105–115, 2014.
- [265] A Meola. What is the Internet of Things (IoT)? Meaning and definition. Business Insider, <http://uk.businessinsider.com/internet-of-things-definition>, 2018. [Online; Accessed 15-Aug-2018].
- [266] M J Metzger. Effects of site, vendor, and consumer characteristics on web site trust and disclosure. *Communication Research*, 33(3):155–179, 2006.
- [267] S Milgram. Behavioral study of obedience. *Journal of Abnormal and Social Psychology*, 67(4):371–378, 1963.
- [268] D Miorandi, S Sicari, F De Pellegrini, and I Chlamtac. Internet of Things: Vision, applications and research challenges. *Ad Hoc Networks*, 10(7):1497–1516, 2012.
- [269] T Moores. Do consumers understand the role of privacy seals in e-commerce? *Communications of the ACM*, 48(3):86–91, 2005.
- [270] Morar Consulting. The dangers of our digital lives. <https://slidex.tips/download/the-dangers-of-our-digital-lives>, 2016. [Online; Accessed 15-Aug-2018].

- [271] E R Morrissey. Sources of error in the coding of questionnaire data. *Sociological Methods & Research*, 3(2):209–232, 1974.
- [272] L F Motiwalla and X B Li. Unveiling consumers’ privacy paradox behaviour in an economic exchange. *International Journal of Business Information Systems*, 23(3):307–329, 2016.
- [273] D Moye. Amazon admits Alexa device eavesdropped on portland family. Huffington Post, https://www.huffingtonpost.co.uk/entry/alexa-eavesdropping-portland-familiy_us_5b0727cae4b0fdb2aa51b23e, 2018. [Online; Accessed 15-Aug-2018].
- [274] S Myagmar, A J Lee, and W Yurcik. Threat modeling as a basis for security requirements. In *Proceedings on the Symposium on Requirements Engineering for Information Security*, 2005.
- [275] E Namey, G Guest, L Thairu, and L Johnson. Data reduction techniques for large qualitative data sets. In *Handbook for Team-Based Qualitative Research*, pages 137–161. 2008.
- [276] P Newman. The Internet of Things 2018 report: How the IoT is evolving to reach the mainstream with businesses and consumers. Business Insider Intelligence, <http://uk.businessinsider.com/the-internet-of-things-2017-report-2018-2-26-1>, 2018. [Online; Accessed 15-Aug-2018].
- [277] R S Nickerson. Null hypothesis significance testing: A review of an old and continuing controversy. *Psychological Methods*, 5(2):241–301, 2000.
- [278] Helen Nissenbaum. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, 2009.
- [279] P A Norberg, D R Horne, and D A Horne. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1):100–126, 2007.
- [280] G Norman. Likert scales, levels of measurement and the “laws” of statistics. *Advances in Health Sciences Education*, 15(5):625–632, 2010.
- [281] P Norman and M Conner. The role of social cognition models in predicting health behaviours: Future directions. In *Predicting Health Behaviour: Research*

- and Practice with Social Cognition Models*, pages 197–225. Open University Press, 1996.
- [282] NPD Connected Intelligence. Consumers and wearables report. <https://www.npd.com/wps/portal/npd/us/news/press-releases/2015/the-demographic-divide-fitness-trackers-and-smartwatches-attracting-very-different-segments-of-the-market-according-to-the-npd-group/>, 2014. [Online; Accessed 15-Aug-2018].
- [283] J R C Nurse, A Erola, I Agrafiotis, M Goldsmith, and S Creese. Smart insiders: Exploring the threat from insiders using the Internet-of-Things. In *Proceedings of ESORICS 2015*, pages 5–14, 2015.
- [284] M C Oetzel and T Gonja. The online privacy paradox: A social representations perspective. In *Proceedings of CHI'11 Extended Abstracts on Human Factors in Computing Systems*, pages 2107–2112, 2011.
- [285] R Opdenakker. Advantages and disadvantages of four interview techniques in qualitative research. *Forum: Qualitative Social Research*, 7(4), 2006.
- [286] D M Oppenheimer, T Meyvis, and N Davidenko. Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, 45(4):867–872, 2009.
- [287] M T Orne. On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist*, 17(11):776–783, 1962.
- [288] K G Orphanides. The Android privacy and security settings you need to know about. Wired, <https://www.wired.co.uk/article/android-privacy-settings-oreo-security>, 2018. [Online; Accessed 15-Aug-2018].
- [289] L Palen and P Dourish. Unpacking privacy for a networked world. In *Proceedings of the 2003 SIGCHI Conference on Human Factors in Computing Systems*, pages 129–136, 2003.
- [290] J Park, Y Ju and J H Ahn. Are people really concerned about their privacy?: Privacy paradox in mobile environment. In *Proceedings of the 15th International Conference on Electronic Business*, pages 123–128, 2015.

- [291] Y J Park. Digital literacy and privacy behavior online. *Communication Research*, 40(2):215–236, 2011.
- [292] S Parkinson, V Eatough, J Holmes, E Stapley, and N Midgley. Framework analysis: A worked example of a study exploring young people’s experiences of depression. *Qualitative Research in Psychology*, 13(2):109–129, 2016.
- [293] M Q Patton. Enhancing the quality and credibility of qualitative analysis. *Health Services Research*, 34(5):1189–1208, 1999.
- [294] K Pearson. Mathematical contributions to the theory of evolution. III. Regression, heredity, and panmixia. *Philosophical Transactions of the Royal Society of London. Series A*, pages 253–318, 1896.
- [295] K Pearson. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 50(302):157–175, 1900.
- [296] E Peer, L Brandimarte, S Samat, and A Acquisti. Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70:153–163, 2017.
- [297] E Peer, J Vosgerau, and A Acquisti. Reputation as a sufficient condition for data quality on Amazon Mechanical Turk. *Behavior Research Methods*, 46(4):1023–1031, 2014.
- [298] I Pentina, L Zhang, H Bata, and Y Chen. Exploring privacy paradox in information-sensitive mobile app adoption: A cross-cultural comparison. *Computers in Human Behavior*, 65:409–419, 2016.
- [299] C Perera, C McCormick, A K Bandara, and B A Price. Privacy-by-Design framework for assessing Internet of Things applications and platforms. In *Proceedings of the 6th International Conference on the Internet of Things*, pages 83–92, 2016.
- [300] C Perera, R Ranjan, L Wang, S U Khan, and A Y Zomaya. Big data privacy in the Internet of Things era. *IT Professional*, 17(3):32–39, 2015.

- [301] C Perera, A Zaslavsky, P Christen, and D Georgakopoulos. Context aware computing for the Internet of Things: A survey. *IEEE Communications Surveys and Tutorials*, 16(1):414–454, 2014.
- [302] A Perez and S Zeadally. Privacy issues and solutions for consumer wearables. *IT Professional*, 20(4):46–56, 2017.
- [303] C Phelan, C Lampe, and P Resnick. It’s creepy, but it doesn’t bother me. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 5240–5251, 2016.
- [304] S Pike, M Kellely, and A Gelnaw. Measuring U.S. privacy sentiment: An IDC special report. <http://www.idc.com/getdoc.jsp?containerId=prUS42253017>, 2017. [Online; Accessed 15-Aug-2018].
- [305] E M Pilke. Flow experiences in information technology use. *International Journal of Human-Computer Studies*, 61(3):347–357, 2004.
- [306] S Pizza, B Brown, D McMillan, and A Lampinen. Smartwatch in vivo. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 5456–5469, 2016.
- [307] C Pope, S Ziebland, and N Mays. Analysing qualitative data. *British Medical Journal*, 320(7227):114–116, 2000.
- [308] M E Porter and J E Heppelmann. How smart, connected products are transforming competition. *Harvard Business Review*, 92(11):64–88, 2014.
- [309] S Pötzsch. Privacy awareness: A means to solve the privacy paradox? *The Future of Identity in the Information Society in IFIP Advances in Information and Communication Technology*, 298:226–236, 2009.
- [310] S Preibusch. Guide to measuring privacy concern: Review of survey and observational instruments. *International Journal of Human-Computer Studies*, 71(12):1133–1143, 2013.
- [311] M Presser. The rise of IoT - Why today? IEEE Internet of Things Newsletter, <https://iot.ieee.org/newsletter/january-2016/the-rise-of-iot-why-today.html>, 2016. [Online; Accessed 15-Aug-2018].

- [312] C N Quinn. *Engaging learning: Designing e-learning simulation games*. John Wiley & Sons, 2005.
- [313] K Quinn. Why we share: A uses and gratifications approach to privacy regulation in social media use. *Journal of Broadcasting & Electronic Media*, 60(1):61–86, 2016.
- [314] N H A Rahman. Privacy disclosure risk: Smartphone user guide. *International Journal of Mobile Network Design and Innovation*, 5(1):2–8, 2013.
- [315] P Rajivan and J Camp. Influence of privacy attitude and privacy cue framing on android app choices. In *Proceedings of the 12th Symposium on Usable Privacy and Security*, 2016.
- [316] R Raman, A Lal, and K Achuthan. Serious games based approach to cyber security concept learning: Indian context. In *Proceedings of the 2014 International Conference on Green Computing Communication and Electrical Engineering*, 2014.
- [317] K Renaud, M Volkamer, and A Renkema-Padmos. Why doesn't Jane protect her privacy? *Privacy Enhancing Technologies in Lecture Notes in Computer Science*, 8555:244–262, 2014.
- [318] K Renaud and V Zimmermann. Ethical guidelines for nudging in information security & privacy. *International Journal of Human-Computer Studies*, 120:22–35, 2018.
- [319] F Rheingans, B Cikit, and C P H Ernst. The potential influence of privacy risk on activity tracker usage: A study. In *The Drivers of Wearable Device Usage*, pages 25–35. Springer, 2016.
- [320] J Ricci, I Baggili, and F Breitingner. Watch what you wear: Smartwatches and sluggish security. In *Wearable Technologies: Concepts, Methodologies, Tools, and Applications*, pages 1458–1478. IGI Global, 2018.
- [321] C Richards, C W Thompson, and N Graham. Beyond designing for motivation: The importance of context in gamification. In *Proceedings of the 1st SIGCHI Annual Symposium on Computer-Human Interaction in Play*, pages 217–226, 2014.

- [322] J Ritchie and L Spencer. Qualitative data analysis for applied policy research. In *Analyzing qualitative data*, pages 173–194. Routledge, 1994.
- [323] J Ritchie, L Spencer, and W O’Connor. Carrying out qualitative analysis. In *Qualitative Research Practice: A Guide for Social Science Students and Researchers*, pages 219–262. Sage Publications, 2003.
- [324] U Ritterfeld, M Cody, and P Vorderer. *Serious games: Mechanisms and effects*. Routledge, 2009.
- [325] J Rivera and R van der Meulen. Gartner says the Internet of Things installed base will grow to 26 billion units by 2020. Gartner, <http://www.gartner.com/newsroom/id/2636073>, 2013. [Online; Accessed 15-Aug-2018].
- [326] A H Robertson. The United Nations Covenant on Civil and Political Rights and the European Convention on Human Rights. *British Yearbook International Law*, 43:21, 1968.
- [327] B Robinson. With a different Marx: Value and the contradictions of Web 2.0 capitalism. *The Information Society*, 31(1):44–51, 2015.
- [328] R W Rogers. Cognitive and psychological processes in fear appeals and attitude change: A revised theory of protection motivation. In *Social Psychophysiology: A Sourcebook*, pages 153–176. Guilford Publications, 1983.
- [329] M Rouse. IoT devices. TechTarget, <https://internetofthingsagenda.techtarget.com/definition/IoT-device>, 2018. [Online; Accessed 15-Aug-2018].
- [330] J Rowley. Conducting research interviews. *Management Research Review*, 35(3/4):260–271, 2012.
- [331] R M Ryan and E L Deci. Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary Educational Psychology*, 25(1):54–67, 2000.
- [332] M A Sasse, D Ashenden, D Lawrence, L Coles-Kemp, I Fléchais, and P Kearney. Human vulnerabilities in security systems. Human Factors Working Group White Paper, <https://pdfs.semanticscholar.org/38b4/36a07f78056a82df1e9228b87ca145f09f9c.pdf>, 2007. [Online; Accessed 15-Aug-2018].

- [333] S Sawilowsky. New effect size rules of thumb. *Journal of Modern Applied Statistical Methods*, 8(2):467–474, 2009.
- [334] B Schneier. *Data and Goliath*. W. W. Norton & Company, 2015.
- [335] m c schraefel, R Gomer, A Alan, E Gerding, and C Maple. The Internet of Things: Interaction challenges to meaningful consent at scale. *Interactions*, 24(6):26–33, 2017.
- [336] P M Schwartz. The EU-US privacy collision: A turn to institutions and procedures. *Harvard Law Review*, 126(7):1966–2009, 2013.
- [337] D Seifert. Secret program gives NSA, FBI backdoor access to Apple, Google, Facebook, Microsoft data. The Verge, <https://www.theverge.com/2013/6/6/4403868/nsa-fbi-mine-data-apple-google-facebook-microsoft-others-prism>, 2013. [Online; Accessed 15-Aug-2018].
- [338] P Shannon-Baker. Making paradigms meaningful in mixed methods research. *Journal of Mixed Methods Research*, 10(4):319–334, 2016.
- [339] R Shay, S Komanduri, A L Durity, P S Huh, M L Mazurek, S M Segreti, B Ur, L Bauer, N Christin, and L F Cranor. Can long passwords be secure and usable? In *Proceedings of the 2014 SIGCHI Conference on Human Factors in Computing Systems*, pages 2927–2936, 2014.
- [340] S Sheng, M Holbrook, P Kumaraguru, L F Cranor, and J Downs. Who falls for phish?: A demographic analysis of phishing susceptibility and effectiveness of interventions. In *Proceedings of the 2010 SIGCHI Conference on Human Factors in Computing Systems*, pages 373–382, 2010.
- [341] S Sheng, B Magnien, and P Kumaraguru. Anti-phishing Phil: The design and evaluation of a game that teaches people not to fall for phish. In *Proceedings of the 3rd Symposium on Usable Privacy and Security*, pages 88–99, 2007.
- [342] M Sherif, D Taub, and C I Hovland. Assimilation and contrast effects of anchoring stimuli on judgments. *Journal of Experimental Psychology*, 55(2):150–155, 1958.
- [343] F Shih, I Liccardi, and D Weitzner. Privacy tipping points in smartphones privacy preferences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 807–816, 2015.

- [344] M Shuttleworth and L T Wilson. Experimental research. <https://explorable.com/scientific-control-group>, 2010. [Online; Accessed 15-Aug-2018].
- [345] S Sicari, A Rizzardi, L A Grieco, and A Coen-Porisini. Security, privacy and trust in Internet of Things: The road ahead. *Computer Networks*, 76:146–164, 2015.
- [346] D Silverman. *Interpreting qualitative data*. Sage, 2015.
- [347] H A Simon. Theories of bounded rationality. *Decision and Organization*, 1(1):161–176, 1972.
- [348] M Siponen, S Pahlila, and M A Mahmood. Compliance with information security policies: An empirical investigation. *Computer*, 43(2):64–71, 2010.
- [349] M T Siponen. A conceptual foundation for organizational information security awareness. *Information Management & Computer Security*, 8(1):31–41, 2000.
- [350] P Slovic. Perception of risk. *Science*, 236(4799):280–285, 1987.
- [351] D P Snyman, H Kruger, and W D Kearney. I shall, we shall, and all others will: Paradoxical information security behaviour. *Information & Computer Security*, 26(3):290–305, 2018.
- [352] D J Solove. *Understanding privacy*. Harvard University Press, 2008.
- [353] D J Solove. *Nothing to hide: The false tradeoff between privacy and security*. Yale University Press, 2011.
- [354] C Spearman. The proof and measurement of association between two things. *The American Journal of Psychology*, 15(1):72–101, 1904.
- [355] S Spiekermann. The challenges of privacy by design. *Communications of the ACM*, 55(7):38–40, 2012.
- [356] S Spiekermann, J Grossklags, and B Berendt. E-privacy in 2nd generation E-commerce: Privacy preferences versus actual behavior. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, pages 38–47, 2001.
- [357] B Stanton, M F Theofanos, S S Prettyman, and S Furman. Security fatigue. *IT Professional*, 18(5):26–32, 2016.

- [358] S S Stevens. On the theory of scales of measurement. *Science*, 103(2684):677–680, 1946.
- [359] C E Stout. Using psychology to counter terrorism at the personal and community level. In *Psychology of Terrorism: Coping with the Continued Threat*, pages 1–32. Praeger, 2004.
- [360] F Strack. “Order effects” in survey research: Activation and information functions of preceding questions. In *Context Effects in Social and Psychological Research*, pages 23–34. Springer, 1992.
- [361] Student. The probable error of a mean. *Biometrika*, 6(1):1–25, 1908.
- [362] F Stutzman, R Gross, and A Acquisti. Silent listeners: The evolution of privacy and disclosure on Facebook. *Journal of Privacy and Confidentiality*, 4(2):7–41, 2013.
- [363] A Suknot, T Chavez, N Rackley, and P G Kelley. Immaculacy: A game of privacy. In *Proceedings of the 1st ACM SIGCHI Annual Symposium on Computer-Human Interaction in Play*, pages 383–386, 2014.
- [364] M Swan. Sensor mania! The Internet of Things, wearable computing, objective metrics, and the quantified self 2.0. *Journal of Sensor and Actuator Networks*, 1(3):217–253, 2012.
- [365] P Syverson. The paradoxical value of privacy. In *Proceedings of the 2nd Annual Workshop on Economics and Information Security*, 2003.
- [366] R Taylor. Interpretation of the correlation coefficient: A basic review. *Journal of Diagnostic Medical Sonography*, 6(1):35–39, 1990.
- [367] R H Thaler and C R Sunstein. *Nudge: Improving decisions about health, wealth and happiness*. Yale University Press, 2008.
- [368] The Economist. Where the smart is. <https://www.economist.com/news/business/21700380-connected-homes-will-take-longer-materialise-expected-where-smart>, 2016. [Online; Accessed 15-Aug-2018].
- [369] S Trepte, T Dienlin, and L Reinecke. Risky behaviors: How online experiences influence privacy behaviors. In *From the Gutenberg Galaxy to the Google Galaxy. Surveying Old and New Frontiers after 50 Years of DGPK*, pages 225–244. UVK, 2014.

- [370] A Tversky and D Kahneman. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2):207–232, 1973.
- [371] E S Udoh and A Alkharashi. Privacy risk awareness and the behavior of smart-watch users: A case study of Indiana University students. In *Proceedings of the 1st Future Technologies Conference*, pages 926–931, 2016.
- [372] A Ukil, J Sen, and S Koilakonda. Embedded security for Internet of Things. In *Proceedings of the 2nd National Conference on Emerging Trends and Applications in Computer Science*, pages 1–6, 2011.
- [373] United States. The Health Insurance Portability and Accountability Act (HIPAA), 2004.
- [374] A Uskov and B Sekar. Smart gamification and smart serious games. In *Fusion of Smart, Multimedia and Computer Gaming Technologies*, pages 7–36. Springer, 2015.
- [375] J Vaidya, D Lorenzi, B Shafiq, S Chun, and N Badar. Teaching privacy in an interactive fashion. In *Proceedings of the 2014 Information Security Curriculum Development Conference*, 2014.
- [376] T L M Van Kasteren, G Englebienne, and B J Kröse. Activity recognition using semi-markov models on real world smart home datasets. *Journal of Ambient Intelligence and Smart Environments*, 2(3):311–325, 2010.
- [377] J P Vasseur and A Dunkels. *Interconnecting smart objects with IP: The next Internet*. Morgan Kaufmann, 2010.
- [378] G A Veltri and A Ivchenko. The impact of different forms of cognitive scarcity on online privacy disclosure. *Computers in Human Behavior*, 73:238–246, 2017.
- [379] B Voyer. ‘Nudging’ behaviours in healthcare: Insights from behavioural economics. *British Journal of Healthcare Management*, 21(3):130–135, 2015.
- [380] S Wall and F Healy. Usability testing of smarter heating controls. UK Department for Energy and Climate Change, https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/266220/usability_testing_smarter_heating_controls.pdf, 2013. [Online; Accessed 15-Aug-2018].

- [381] H Walsh. Should I buy a fitness tracker or a smartwatch? Which?, <https://www.which.co.uk/reviews/fitness-watches-and-activity-trackers/article/should-i-buy-a-fitness-tracker-or-a-smartwatch>, 2017. [Online; Accessed 15-Aug-2018].
- [382] G Walsham. The emergence of interpretivism in IS research. *Information Systems Research*, 6(4):376–394, 1995.
- [383] N Wang, B Zhang, B Liu, and H Jin. Investigating effects of control and ads awareness on Android users’ privacy behaviors and perceptions. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services*, pages 373–382, 2015.
- [384] Y Wang, P G Leon, A Acquisti, L F Cranor, A Forget, and N Sadeh. A field trial of privacy nudges for Facebook. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*, pages 2367–2376, 2014.
- [385] Y Wang, G Norcie, S Komanduri, A Acquisti, P G Leon, and L F Cranor. “I regretted the minute I pressed share”: A qualitative study of regrets on Facebook. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*, page 10, New York, New York, USA, 2011. ACM Press.
- [386] S D Warren and L D Brandeis. The right to privacy. *Harvard Law Review*, 4(5):193–220, 1890.
- [387] J Q Whitman. The two western cultures of privacy: Dignity versus liberty. *Yale Law Journal*, 113(6):1151–1221, 2004.
- [388] A Whitten and J D Tygar. Why Johnny can’t encrypt: A usability evaluation of PGP 5.0. In *Proceedings of the 8th conference on USENIX Security Symposium*, pages 14–29, 1999.
- [389] S Wiedenbeck, J Waters, J C Birget, A Brodskiy, and N Memon. Authentication using graphical passwords: Effects of tolerance and image choice. In *Proceedings of the 1st Symposium on Usable Privacy and Security*, pages 1–12, 2005.
- [390] A Wieneke, C Lehrer, and R Zeder. Privacy-related decision-making in the context of wearable use. In *Proceedings of the 20th Pacific Asia Conference on Information Systems*, 2016.

- [391] F Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83, 1945.
- [392] G J Wilde. The theory of risk homeostasis: Implications for safety and health. *Risk Analysis*, 2(4):209–225, 1982.
- [393] P Wilkinson. A brief history of serious games. In *Entertainment Computing and Serious Games*, pages 17–41. Springer, 2016.
- [394] P Wouters, C Van Nimwegen, H Van Oostendorp, and E D Van Der Spek. A meta-analysis of the cognitive and motivational effects of serious games. *Journal of Educational Psychology*, 105(2):249–265, 2013.
- [395] A C Yeo, M Rahim, and Y Y Ren. Use of persuasive technology to change end user’s IT security aware behavior: A pilot study. *International Journal of Human and Social Sciences*, 4(9):673–679, 2009.
- [396] S Youn. Determinants of online privacy concern and its influence on privacy protection behaviors among young adolescents. *Journal of Consumer Affairs*, 43(3):389–418, 2009.
- [397] A M Zafeiropoulou, D E Millard, C Webber, and K O’Hara. Unpicking the privacy paradox: Can structuration theory help to explain location-based privacy decisions? In *Proceedings of the 5th Annual ACM Web Science Conference*, pages 463–472, 2013.
- [398] P A Zandbergen. Accuracy of iPhone locations: A comparison of assisted GPS, WiFi and cellular positioning. *Transactions in GIS*, 13(1):5–25, 2009.